ALMA MATER STUDIORUM — UNIVERSITÀ DI BOLOGNA

DEIS - Dipartimento di Elettronica, Informatica e Sistemistica
Dottorato in Ingegneria Elettronica, Informatica e Telecomunicazioni

Ciclo XXVII
Settore Concorsuale: 09/H1
Settore Disciplinare: ING-INF/05

ENGINEERING COMPLEX COMPUTATIONAL ECOSYSTEMS

Candidato

Dott. Ing. DANILO PIANINI

Supervisore

Chiar.mo Prof. Ing. MIRKO VIROLI

Tutor

Chiar.mo Prof. Ing. ANTONIO NATALI

Coordinatore

Chiar.mo Prof. Ing. ALESSANDRO VANELLI CORALLI

FINAL EXAMINATION YEAR 2015

Considerate la vostra semenza:
fatti non foste a viver come bruti,
ma per seguir virtute e canoscenza.

Acknowledgements

Most of the material that got integrated as a single work in this thesis has been written in publications of which I am co-author. I am thankful to those great researchers I have collaborated with, for their suggestions, their guidance and their help. I'd like to thank first Mirko Viroli, who supervised me for the past three years and tried to make me a scientist. Sara Montagna, who shares the APICe laboratory with me, also deserves a special mention. I want to thank the work packages leaders of SAPERE, a European project that kickstarted me in the world of research: Franco Zambonelli, Giovanna di Marzo, Simon Dobson, Marco Mamei, Alois Ferscha (who also supervised me for three months in Linz) and Andrea Omicini. A special thank to Jacob "Jake" Beal, that supervised (and also hosted) me during my last three months in the United States. Finally, I desire to thank these other great researchers and professionals I have collaborated with: Bernhard Anzengruber, Jose-Luis Fernandez Marquez, Stefano Mariani, Ronaldo Menezes, Jussi Nieminen, Matteo Risoldi, Stefano Sebastio, Graeme Stevenson, Andrea Vandin, Sascia Virruso, Juan "Erica" Ye.

Danilo Pianini, 13th March 2015

Contents

Abstract

Self-organising pervasive ecosystems of devices are set to become a major vehicle for delivering infrastructure and end-user services. The inherent complexity of such systems poses new challenges to those who want to dominate it by applying the principles of engineering.

The recent growth in number and distribution of devices with decent computational and communicational abilities, that suddenly accelerated with the massive diffusion of smartphones and tablets, is delivering a world with a much higher density of devices in space. Also, communication technologies seem to be focussing on short-range device-to-device (P2P) interactions, with technologies such as Bluetooth and Near-Field Communication gaining greater adoption.

Locality and situatedness become key to providing the best possible experience to users, and the classic model of a centralised, enormously powerful server gathering and processing data becomes less and less efficient with device density. Accomplishing complex global tasks without a centralised controller responsible of aggregating data, however, is a challenging task. In particular, there is a local-to-global issue that makes the application of engineering principles challenging at least: designing device-local programs that, through interaction, guarantee a certain global service level.

In this thesis, we first analyse the state of the art in coordination systems, then motivate the work by describing the main issues of pre-existing tools and practices and identifying the improvements that would benefit the design of such complex software ecosystems.

The contribution can be divided in three main branches. First, we introduce a novel simulation toolchain for pervasive ecosystems, designed for allowing good expressiveness still retaining high performance. Second, we leverage existing coordination models and patterns in order to create new spatial structures. Third, we introduce a novel language, based on the existing "Field Calculus" and integrated with the aforementioned toolchain, designed to be usable for practical aggregate programming.

1 About this thesis

This chapter briefly introduces the content of this work, the main problems tackled, the reasons why it is relevant, and how the work is organised. Also, a list of publications that directly contribute to this thesis is provided.

This thesis contributes to the emergent field of computational ecosystems, namely complex systems in which multiple distributed services and applications live and evolve, possibly interacting among them. Also, the situation is complicated by the fact that in general there is no guarantee of a centralised network node or service that may provide the computation required, as it is the case in cloud systems.

Such complex systems require novel engineering procedures and tools, because they deeply differ from the systems classical software engineering is devoted to. This thesis contributes to the engineering of complex computational ecosystems in four areas:

Language area Evolution of computer science is, in a sense, evolution of paradigms and languages. One of the possible approaches to the novel field subject of this research is to work at this fundamental level, providing proper abstractions that provide the desired expressive power leaving out unnecessary low-level details.

Toolchain area In classic systems, powerful testing and verification tools exist to that guarantee certain properties to hold. Testing happen on the very same code that is going to be deployed, and it is very often applicable at the unit level. A similar tool chain is not effective in testing complex ecosystems: the main reason is that most of the desired behaviours have little sense for a single device, but rather are a property of the aggregate of devices itself. Moreover, the environment in which the system is deployed may consistently influence the behaviour. In such complexity, simulation is very often the only way out.

Design pattern area Classical design patterns can not be applied for the engineering of distributed ecosystems. However, a different kind of patterns can be found in literature: a library of local (device level) behaviours whose interaction result in global behaviours is available. This thesis relies on existing work and extends it with a "channel" pattern and, more notably, with an anticipative adaptation pattern.

Application prototype area Which applications can we build upon this new paradigm? Are

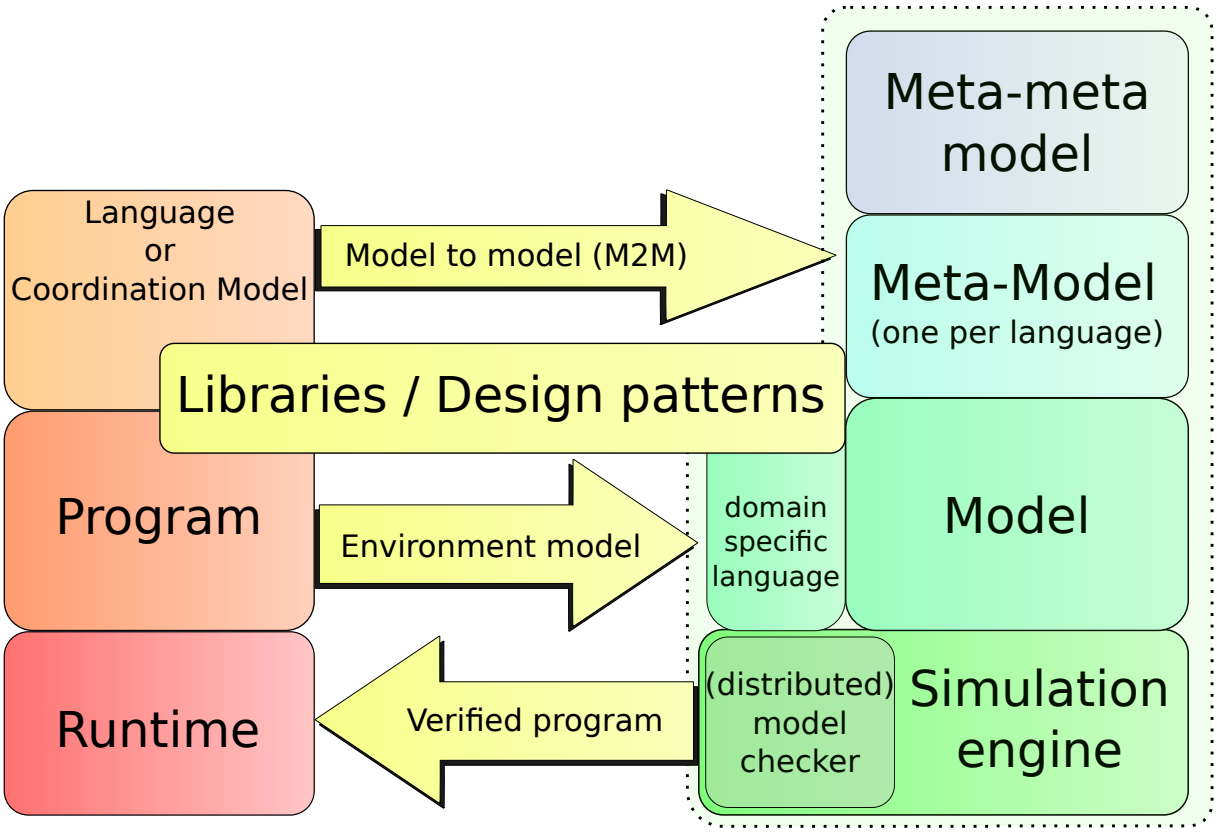

Figure 1.1: Abstract structure of the engineering process and toolchain for complex ecosystems

the presented languages and patterns really helpful in satisfying the requirements? In this work, some insights are provided for each technique, tool, or language, in order to put the proposed solutions at work.

In Figure 1.1 we give an overview of the engineering process behind a computational ecosystem, also showing which components would be required in order to have a proper toolchain. The goal is to have one or more distributed *programs* that run on a *runtime*. Such programs are written in a specific *language*, or rely on some *coordination model*, and possibly use *libraries or design patterns* (in the sense of spatial pattern, as discussed above). Designing such programs and testing them prior to deployment requires another series of tools, which are enclosed in a dashed shape. In the schema, we suppose a meta-meta model, which is instanced once for each language or coordination model that should be supported by the development and testing toolchain, generating a meta-model. In order to test some program, it must be fed to the meta-model along with a model of the environment, since it has a strong influence on the behaviour of the whole system. The environment description could be possibly provided in a dedicated domain specific language, that, along with the program to simulate, constitutes the model. Such model can be tested and simulated, provided that a suitable engine exists. Once verified, the

program can be deployed on the actual runtime.

1.1 Structure of this thesis

This thesis is organised in five parts.

Part I introduces the work, describing the background and motivating the thesis. It presents the evolution and the current state of the pervasive communication technologies, which are the practical devices on which computational ecosystems are designed to run. Also, it discusses the state of the art in the software world for what concerns computational ecosystems. It describes both existing middlewares and engineering tools.

Part II presents the first contribution of this thesis, namely a toolchain for engineering computational ecosystems. With respect to Figure 1.1, this part describes a novel work that occupies the whole structure on the right-hand side of the schema (enclosed with a dashed line).

Part III focuses on the left-hand side of Figure 1.1, introducing two novel spatial patterns (which fit in the yellow box), a process algebra (that belongs to the "Coordination model" balloon), and a series of case studies demonstrating the effectiveness of pervasive ecosystems (which would be best inserted in the orange "Program" box).

Part IV focusses on a more fundamental approach, trying to define novel languages suitable for programming aggregates of devices. Most of the content of this part fits into the "Language" top left box in Figure 1.1. One of the presented languages also features a meta-model for the toolchain mentioned in Part II and a runtime, that make it a practical possibility for programming aggregates.

Part V draws conclusions and paves the way for future works.

1.2 List of publications

- Danilo Pianini, Sascia Virruso, Ronaldo Menezes, Andrea Omicini, and Mirko Viroli. Self organization in coordination systems using a wordnet-based ontology. In *SASO*, pages 114–123, 2010

- Danilo Pianini, Mirko Viroli, and Sara Montagna. A simulation framework for pervasive services ecosystems. In *Proceedings of the 12th Workshop on Objects and Agents, Rende (CS), Italy, Jul 4-6, 2011*, pages 150–157, 2011

- Sara Montagna, Mirko Viroli, Matteo Risoldi, Danilo Pianini, and Giovanna Di Marzo Serugendo. Self-organising pervasive ecosystems: A crowd evacuation example. In *Software Engineering for Resilient Systems - Third International Workshop, SERENE 2011, Geneva, Switzerland, September 29-30, 2011. Proceedings*, pages 115–129, 2011

- Danilo Pianini, Sara Montagna, and Mirko Viroli. A chemical inspired simulation framework for pervasive services ecosystems. In *Proceedings of the Federated Conference on Computer Science and Information Systems*, pages 667–674. IEEE Computer Society Press, 2011

- Sara Montagna, Danilo Pianini, and Mirko Viroli. Gradient-based self-organisation patterns of anticipative adaptation. In *Proceedings of 6th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2012)*, pages 169–174, September 2012

- Sara Montagna, Danilo Pianini, and Mirko Viroli. A model for drosophila melanogaster development from a single cell to stripe pattern formation. In *Proceedings of the ACM Symposium on Applied Computing, SAC 2012, Riva, Trento, Italy, March 26-30, 2012*, pages 1406–1412, 2012

- Sara Montagna, Mirko Viroli, Danilo Pianini, and Jose Luis Fernandez-Marquez. Towards a comprehensive approach to spontaneous self-composition in pervasive ecosystems. In Flavio De Paoli and Giuseppe Vizzari, editors, *Proceedings of the 13th Workshop on Objects and Agents*. CEUR-WS, 12 September 2012

- Mirko Viroli, Danilo Pianini, Sara Montagna, and Graeme Stevenson. Pervasive ecosystems: a coordination model based on semantic chemistry. In Sascha Ossowski, Paola Lecca, Chih-Cheng Hung, and Jiman Hong, editors, *27th Annual ACM Symposium on Applied Computing (SAC 2012)*, Riva del Garda, TN, Italy, 26–30 March 2012. ACM

- Mirko Viroli, Danilo Pianini, and Jacob Beal. Linda in space-time: an adaptive coordination model for mobile ad-hoc environments. In Marjan Sirjani, editor, *Coordination Languages and Models*, volume 7274 of *LNCS*, pages 212–229. Springer-Verlag, June 2012. Proceedings of the 14th Conference of Coordination Models and Languages (Coordination 2012), Stockholm (Sweden), 14-15 June

- Bernhard Anzengruber, Danilo Pianini, Jussi Nieminen, and Alois Ferscha. Predicting social density in mass events to prevent crowd disasters. In Adam Jatowt, Ee-Peng Lim, Ying Ding, Asako Miura, Taro Tezuka, Gaël Dias, Katsumi Tanaka, Andrew J. Flanagin, and Bing Tian Dai, editors, *SocInfo*, volume 8238 of *Lecture Notes in Computer Science*, pages 206–215. Springer, 2013

- Graeme Stevenson, Juan Ye, Simon Dobson, Danilo Pianini, Sara Montagna, and Mirko Viroli. Combining self-organisation, context-awareness and semantic reasoning: the case of

resource discovery in opportunistic networks. In Sung Y. Shin and José Carlos Maldonado, editors, *SAC*, pages 1369–1376, Coimbra, Portugal, 18–22 March 2013. ACM

- Danilo Pianini, Sara Montagna, and Mirko Viroli. Chemical-oriented simulation of computational systems with Alchemist. *Journal of Simulation*, 2013
- Danilo Pianini, Mirko Viroli, Franco Zambonelli, and Alois Ferscha. HPC from a self-organisation perspective: The case of crowd steering at the urban scale. In *International Conference on High Performance Computing & Simulation, HPCS 2014, Bologna, Italy, 21-25 July, 2014*, pages 460–467, 2014
- Danilo Pianini, Stefano Sebastio, and Andrea Vandin. Distributed statistical analysis of complex systems modeled through a chemical metaphor. In *International Conference on High Performance Computing & Simulation, HPCS 2014, Bologna, Italy, 21-25 July, 2014*, pages 416–423, 2014
- Franco Zambonelli, Andrea Omicini, Bernhard Anzengruber, Gabriella Castelli, Francesco L. DeAngelis, Giovanna Di Marzo Serugendo, Simon Dobson, Jose Luis Fernandez-Marquez, Alois Ferscha, Marco Mamei, Stefano Mariani, Ambra Molesini, Sara Montagna, Jussi Nieminen, Danilo Pianini, Matteo Risoldi, Alberto Rosi, Graeme Stevenson, Mirko Viroli, and Juan Ye. Developing pervasive multi-agent systems with nature-inspired coordination. *Pervasive and Mobile Computing*, 2014. Special Issue on "10 years of Pervasive Computing" in honor of Chatschik Bisdikian
- Danilo Pianini, Jacob Beal, and Mirko Viroli. Practical aggregate programming with PROTELIS. In *ACM Symposium on Applied Computing (SAC 2015)*, 2015. To appear
- Mirko Viroli, Sara Montagna, Danilo Pianini, Graeme Stevenson, and Franco Zambonelli. A coordination model of pervasive service ecosystems. *Science of Computer Programming*, 2014. Submitted for publication
- Jacob Beal, Mirko Viroli, Danilo Pianini, and Ferruccio Damiani. Scale-independent computations in situated networks. *IEEE Transaction on Emergent Topics in Computing*, 2014. Submitted for publication

Part I

Background and Motivation

2
Pervasive computing technologies

It is no mystery that, with the huge progresses of miniaturisation, computational-capable devices are populating the world. The diffusion got momentum with affordable "personal computers" that made their way in a one-device-per-family world. Laptops boosted the process, offering user a personal and mobile device.

The pervasive revolution, however took place when the phones became "smart", and with the subsequent improvements in the communication technologies. In this chapter we try to briefly walk the path of success of personal mobile devices, describing also the probable newcomer, namely wearable devices. We also focus on the current status of the communication protocols, their range and usage, and we try to foresee which world are we going to build if the current trend continues.

2.1 Smart, portable devices

When Apple in 2007 released the first iPhone (Figure 2.1), a mobile revolution started. Even though other manufacturers proposed products similar to iPhone under the point of view of communication technologies and computational capabilities in the same time frame (e.g. Nokia N810), the Apple's smartphone was the first gaining widespread adoption. It is not really relevant for this work to understand if the branding, the design, the multi-touch finger based interaction UI or the feature set was the key of its commercial success: what really matters is that, starting 2007, every person began to carry with her a personal device featuring both the abilities to communicate and compute. The reason, besides the success of iPhone in the higher segment of the phone market, is mainly to attribute to the widespread diffusion of similar but cheaper and less powerful devices in the lower segments. This kind of devices ultimately pushed the market share of the dominating today's mobile operating system, Google's Android.

This trend towards a higher diffusion got another leap forward three years later, in 2010, when a device with a feature set similar to iPhone, but with no phone abilities and a bigger screen was released: the iPad. As iPhone gave new vitality and perspectives to the phone arena, iPad revitalised a market that was languishing: the tablets. Tablet devices as they were conceived before iPad launched were nothing more than small laptops with a screen that could be rotated or detached from the keyboard, and a touch-screen normally used with the help of a pen. They

Figure 2.1: Apple iPhone is one of the first devices without any physical keyboard, and the first smartphone to gain worldwide success. *Source: Wikimedia.*

changed from devices designed for a professional niche to widespread tools, up to the point that Gartner forecasts their shipments to overtake in 2015 those of desktop PC and laptops aggregated [1].

The new frontier of pervasive is probably the wearable technology. Smartphones were precursors, they substituted mobile phones introducing new features and they have potential to substitute our wallets (see, for instance, payments through NFC technology) and keys, becoming the only object we need to carry with us in our pockets. Still there are other accessories which are hard to replace, above all watches and glasses. Yes, smartphone can easily show the current time accurately (and, to be honest, also feature phones had this feature well before 2007), but they require the user to pick them from the pocket, turn them on, and sometimes, depending on privacy settings, also unlock them. The whole operation, takes a much longer time and higher effort with respect to just rotating a wrist. This might be the reason why, despite the explosive expansion of smartphones, the wristwatches market did not declined, as an analysis from MarketWatch points out [2].

There is room for manufacturers to create new portable devices. Sony, starting 2012, has produced a series of smartwatches, such as those in Figure 2.2 that pair with a smartphone and provide quick access to some of its functionalities. The success of such solution is not huge, but despite that many other companies are interested in this market: Samsung, Motorola and Apple presented devices meant to replace the classic wristwatch, a clear sign that this market is in expansion. At the time of writing, the main issues that slow the widespread adoption of such solutions are battery duration and dependence on a smartphone.

Another notable attempt to make a common accessory smarter is Google Glass project, depicted in Figure 2.3. Their goal is to enhance the experience of wearing glasses by attaching a device with a camera, an optical head-mounted display, and the abilities to locate itself and communicate with other devices. Google Glass, at the time of writing, are way to expensive (with

[1] `http://www.gartner.com/newsroom/id/2791017`

[2] `http://www.marketwatch.com/story/the-watchs-time-isnt-up-2013-07-01`

Figure 2.2: Smart watches. From left to right: Sony Smartwatch *(Source: Alex S.H. Lin)*, Samsung Galaxy Gear *(Source: Karlis Dambrans)*, Motorola Moto 360.

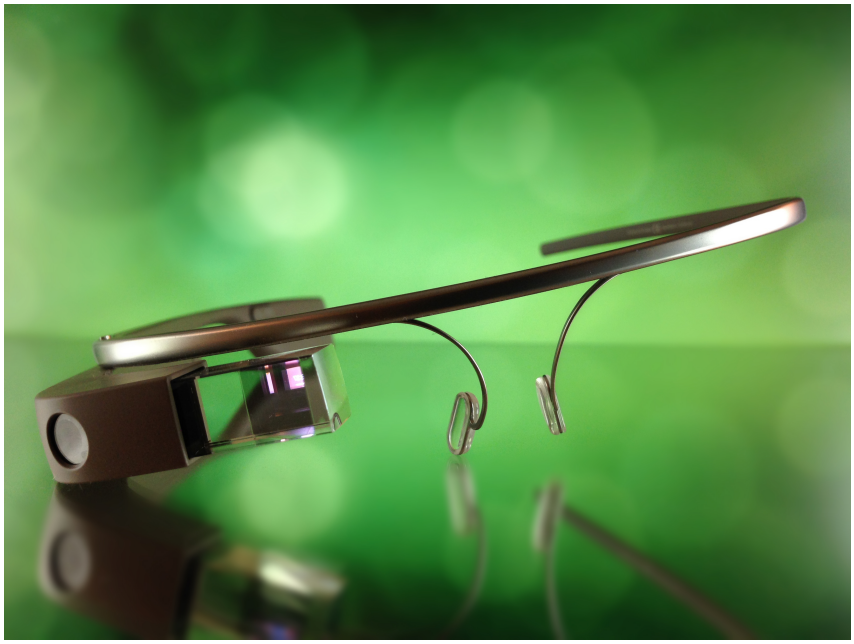

Figure 2.3: Google Glass. In this image, it is possible to see both the camera (on the left hand side) and the semi-transparent head-mounted display. Are those devices going to be part of our everyday life? *Source: Wikimedia.*

the kit sold at $1500) for being able to penetrate the general public, but they are an interesting anticipation of possible future devices.

On the same line of such wearable devices, a discrete success is being achieved by the so called "fit bands". They are bracelets equipped with low energy sensors, mainly accelerometers and gyroscopes, which are used to keep track of user's activity. Depending on the model, they can be used to monitor some user's health parameters, such as the number of steps walked per day or heartbeats. They normally work along with another device, a smartphone or a tablet. Such devices, due to their precise market niche and reasonably low price (the Chinese manufacturer Xiaomi recently introduced a low-end wristband at around $15) are having a notable success.

The wearable devices segment also includes less common devices such as "smart shoes" and materials that can be used to make clothing, such as e-textiles. It is a market in expansion, greatly beneficing from recent increases in performance per watt efficiency. If the trend continues, it is likely that we will more and more powerful wearable devices on sale at cheaper and cheaper prices, and a consequent widespread diffusion. The same sort may occur to other parts of our life: kitchen gear, indoor lights and many other objects are getting more and more "smart" around us. We may, literally, end up with a world where every single object embeds computational and communicational abilities.

A problem arise: how can software engineers deal with such a complexity?

2.2 Communication technologies

Besides miniaturisation, and as a consequence the increase of computational density in space, another factor played a fundamental role in the world of pervasive computing: the ability to communicate, and in particular the ability to rely on wireless communication, which is of paramount importance when considering mobility.

In later years, many communication means arose. They largely differ in terms of range, protocols, and availability. In this section we try to resume the most diffused technologies available on today's devices, but the reader is warned: keep in mind that such technologies are evolving very quickly, and the scenario is incredibly fluid.

2.2.1 International Mobile Telecommunications

This first mean of communication is designed to allow mobile devices to access the Internet from anywhere in the world, relying on the existing mobile phone infrastructure. Such technologies are meant to be used with the standard IP protocols, and they are normally used to get access to public services, in particular to the world wide web. They are not designed for a local peer-to-peer (P2P) communication, and as a consequence they provide no mean to exploit locality. The diffusion of such communication protocols is widespread, in particular among smartphones. Due to the fact that they rely on the mobile phone network, they require a contract with a

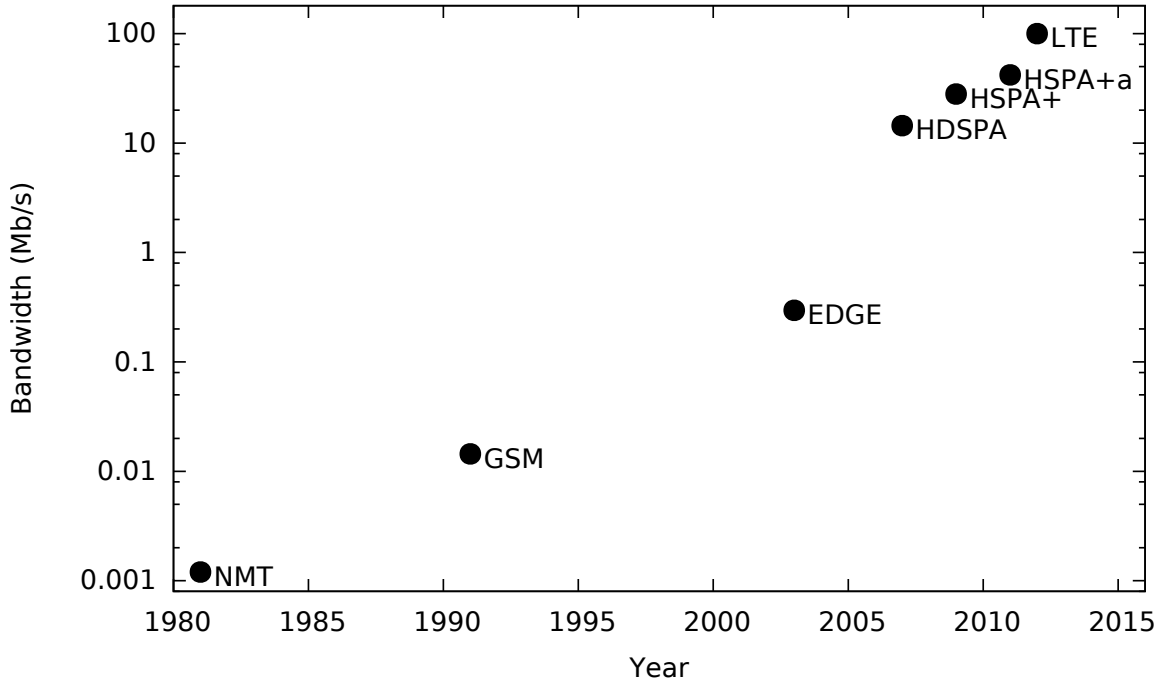

Figure 2.4: Maximum download bandwidth available for mobile devices with time. Each point is labelled with the specific communication technology name. For the most recent multiple-input-multiple-output (MIMO) technologies, such as LTE, a conservative single input channel was used.

mobile telecommunications provider, and as such they are much less diffused in tablets and other portable devices.

The possibility of accessing the Internet from everywhere is probably one of the key bricks that allowed for the huge success of smartphones in today's world. As Figure 2.4 shows, the bandwidth available grew exponentially with time, to the point that in some countries (e.g. in Italy, at the time of writing) the best available mobile connections offer a higher performance than the best available home connection [3]. Such performance unlock the possibility of fully exploiting the possibilities of the world wide web, including cloud services and fruition of multimedia content.

If bandwidth is not currently an issue for international mobile telecommunications, the situation is well different when it comes to device density. Any of us probably experienced network availability issues when participating crowded events, such as concerts or sport events. The current technology, in fact, makes all the network user share the same physical resources: when the device density is too high, there is simply not enough space in the frequency spectrum

[3] At the time of writing, Telecom Italia Mobile offers mobile connections on LTE with a download bandwidth up to 225Mb/s. Fastweb, the company offering the faster solutions for fiber-to-home connections, goes up to 100Mb/s.

to grant a decent bandwidth to everyone. Future networks (5G, and presumably those that will follow) are focussing toward this issue among others [ASS+11]. In particular, a so-called "spectrum crunch" is expected due to the expected traffic increase (thousand fold over this decade and still growing into the next), that could not be faced simply with the foreseen steady increase of the spectrum allocated for mobile communication, and will require technological advances [AARC+12].

2.2.2 WiFi

WiFi technology is the most diffused technology for wireless local networking. It is widely diffused, integrated in all smartphones and tablets and also in other devices, such as printers, gaming consoles and TVs. WiFi devices communicate on a distance that ranges from 20 to 100 meters, depending on the condition of the wireless medium and on the power of the communication devices. The communication speed between two linked devices ranges from 56Mb/s to 300Mb/s.

WiFi was designed to provide wireless access to a local area network. In the most classic "infrastructure mode", wireless devices get connected to a so called access point, which is responsible to route packets among wireless devices and bridge the wireless local area network to the wired backbone. Multiple access point that share the network name (SSID) may be connected using wired network technologies, and they will appear as a single, bigger access point. It is also possible to drop the wired backbone, but specific access points are required.

Also, some WiFi devices provided "ad-hoc mode", allowing multiple devices to directly communicate without an intermediate access point. This working mode was problematic, mainly due to the fact a standard communication protocol for peer-to-peer WiFi communication was missing. This lack was filled with WiFi Direct, which provides a protocol by which one of the devices that want to communicate directly becomes the access point, allowing for direct communication. The most common usage of such a feature are direct file sharing between devices and connection to peripheral devices such as printers or scanners.

2.2.3 Bluetooth and Bluetooth LE

Bluetooth is a technology designed for building energy efficient personal area networks (PANs). Bluetooth devices are assigned a class which identifies the maximum permitted power and, consequently, the maximum operating range. For the most powerful (and power hungry) devices the communication range can go up to 100m. The communication speed ranges from 1Mb/s of the earliest 1.0 version to the 24 Mb/s of version 3.0 and later.

The most interesting features of Bluetooth are not the bandwidth nor the range (WiFi performs better on both), but rather the simple association process and the low power consumption. Thanks to those features Bluetooth found widespread diffusion as a mean to connect low consumption peripherals, such as headsets. Also, it has diffused in cars, and allows user to use the car's audio system as a speakerphone for making calls or listen to music.

Figure 2.5: A iBeacon.

A technology which is often associated with Bluetooth but that is actually a separated and not compatible protocol is Bluetooth LE. The reason why such technologies get associated is that, since the radio frequency used is the same (2.4GHz), dual mode devices can share a single radio antenna. LE stands for Low Energy, and it is the main difference between the two protocols: at the expense of some bandwidth, Bluetooth LE consistently reduce the amount of energy required. Bluetooth LE applications are particularly interesting, and range from health care to fitness to alerts to proximity sensing.

Proximity can be estimated using the received signal strength indicator (RSSI), and the very low power consumption of Bluetooth LE allowed for the realisation of electronic leashing systems, namely systems where an electronic device is paired to an object and can be used in order to compute the relative position. The applications are, for instance, finding of misplaced, out-of-sight devices (when the electronic device is paired with a movable object) and indoor localisation (if the electronic device is located on a still standing object). Relying on this technology, Apple created iBeacon, namely very small (coin sized, see Figure 2.5) electronic devices consisting basically of a battery and a Bluetooth LE device. iBeacons can be attached to objects, and they send a universally unique identifier (UUID) to enabled smartphones in range. If the smartphone can associate the UUID with a position, it can deduce its location relative to the iBeacon. The low energy feature plays a fundamental role in this kind of applications: beacons whose battery would last few hours would be of little practical use. With current technologies, a beacon device can be powered by a standard, rather cheap battery for several months, up to a couple of years. This technology is probably the prelude to precise indoor localisation.

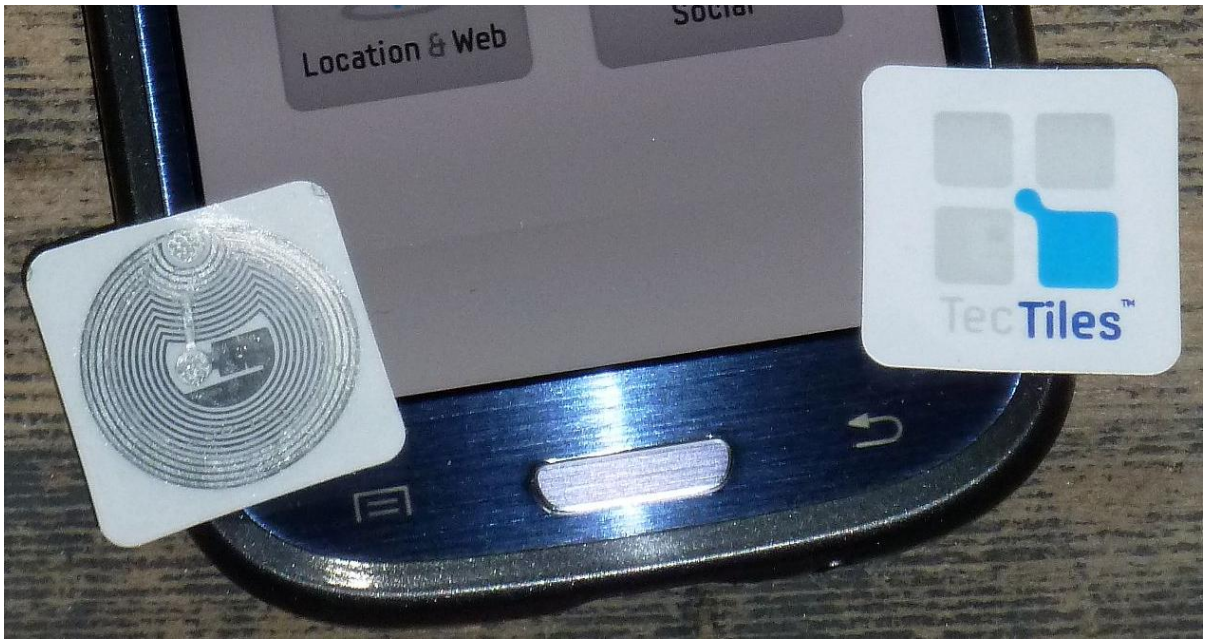

Figure 2.6: Two NFC tags, mounted on stickers. In order to understand the size, they are placed above a Samsung Galaxy S3, a 4-inches smartphone. Moreover, smaller versions of such tags exists, the models pictured here are very common-sized.

2.2.4 NFC

Near Field Communication, or NFC, is a technology designed for low energy communication between two devices in proximity (typically few centimetres). It is designed for low energy consumption rather than high bandwidth: its speed (depending on the specification) ranges from 126kb/s to 242kb/s.

One of the most interesting features of NFC is that one of the two devices (the so-called "tag") can be completely passive, and still carry a small amount of information within (currently between 96 and 4096 bytes). No battery or energy source is required, the information included can be read by active NFC devices in proximity. Having no need of battery at all, NFC tags can really be tiny, even smaller than iBeacons. Figure 2.6 shows two NFC tags: compared with iBeacons, they can be much lighter and thinner. Brought to the world of humans, it is something like creating a sticker with very small sized text: those who have a powerful enough magnifying glass and are close enough to use it properly can read what it says.

The range of applications of such technology is rather broad. The one which was probably most sponsored is contact-less payment, namely the ability to pay just tapping the phone close to the check-out counter. This is a very interesting possibility, and indeed in 2007 there were enthusiastic forecasts [OP07] about its quick diffusion, that did not happen as quick as expected. A number of studies tried to understand the reasons behind this slow adoption, and it appears that reasons are more correlated to marketing and management rather than technological maturity

[PH14, Apa13, TOCH14]. Similarly, if the phone stores identity or access tokens, NFC is a very suitable technology for effectively using such tokens: in this case, its very low range is a nice feature. NFC can also be used as a technology enabler, namely as a mean to securely bootstrap another connection, or join a local network of devices. An example of such use is the Android Beam technology, that relies on NFC in order to establish a Bluetooth connection between two Android devices, transfer a file, then close the Bluetooth connection. Technically, NFC could be used to directly transfer files, but both WiFi and Bluetooth offer much wider bandwidths and range, and as a consequence are preferred for such task. Another interesting application field is mobile device automation: it is achieved by attaching Generally speaking, NFC comes in handy when there is need of a communication mean whose range should be very limited.

2.3 Towards a P2P pervasive continuum?

We are living exciting times. In about five years from the introduction of the technology on the market, almost everybody got a personal smart device always with her. Miniaturisation and power efficiency is constantly growing at a stunning rate, allowing data, communication systems and computation to be spread around in our physical world.

In few decades, we will probably witness the diffusion on computation on everyday object. In such a scenario, the device density will be much higher if compared to the current, up to the point that the aggregation of devices participating the system could be seen as a "pervasive continuum" [ZCF+11]. This continuum is studied under a number of names, including pervasive computing, smart cities, and the Internet of Things; in all cases with the goal of providing a wealth of services in an un-intrusive manner [ZOA+14, CDB+12, Zam12, HDBDM12].

One of the possible strategies is to connect every single device to the Internet, aggregate its information in a remote server, do the necessary computation, then send back eventual results where they are needed. This is the strategy behind cloud computing, which is achieving great success. In particular, we argue, this strategy is interesting when the information could or should be aggregated with information from other, distant sources, or conserved for historical purposes. This path, however, gets harder and harder to follow with device density: besides the obvious increasing on the total information produced, and consequently of the information to transmit and process, there are two other problems: the saturation of the wireless medium, and the locality of information.

Whoever has tried to use its own smartphone in a very crowded environment has probably experienced connection or network issues. The problem, as discussed in Section 2.2.1, is that current technology must share a common medium among all the devices in the same area. Increasing the maximum number of devices per area is one of the goals of the next generation of international mobile telecommunication technologies. One of the proposed approaches is to switch to a very dense array of very small cells by deploying multiple antennas at a very short distance one another, e.g. inside the public illumination poles. Devices nearby the local antenna would connect to it, and the antenna would then connect them to the rest of the network

transparently, in a way somehow similar to the current WiFi "infrastructure mode". Clearly, diffusing such antennas can possibly represent a major infrastructural upgrade, and, potentially, cost.

Another observation is that not every device needs direct access to the Internet to be able to accomplish its task, and this is increasingly true with increased density. Thinking about today's devices, let's consider the current smartwatches and fitbands: the former relies on a smartphone or tablet in order to provide Internet-based services, and the latter uses no Internet connection at all, but just sends data to the smartphone to be processed. Along the line of favouring locality, there is a second advantage which relates to privacy issues: there is no reason to send personal information away, if the system does not need data from distant points nor requires more computational power of the amount available locally. Privacy issues gain great attention recently, especially after the leaks of classified information started in 2010 on Wikileaks and continued with the more recent leaks by Edward Snowden. The content of such documents raised greater attention to privacy issues from general public.

A possible path which would help in both those directions (reduce wireless medium usage and keep data as local as possible) is the usage of local, possibly peer-to-peer interactions. This way of organising communication is already exploited by existing applications.

One notable example is Firechat [4], which got particularly spotlighted during the "Umbrella revolution", namely the sequence of protests that took place in Hong Kong in 2014. Similarly to what was done during Arab springs in 2011 [HDF+11], protesters relied on Internet services to organise and coordinate themselves. The Chinese government policy towards the Internet remains restrictive [Mac08], and services such as Facebook adn Twitter, widely exploited during Arab springs, were already effectively blocked in the land. In short time, other social network were closed (such as Instagram), in order to cut protesters' communication means. At that point, protesters had to find a communication system free of centralisation in order to prevent targeted Internet filters, and Firechat was the answer. Firechat is a messaging application for mobile phones that, when the smartphone has access to the remote Firechat cloud, works as other more famous alternatives, such as WhatsApp[5] and Telegram[6] do. When no access to Firechat servers is available, then the software tries to reach the destination by spreading the message hop-by-hop, building a de-facto mesh network.

Another interesting experiment is Serval Mesh [GSP11]. Serval Mesh accomplishes similar tasks, but it also supports calls and file transfers besides messaging. Its main goal is to provide a networking among users who are in an area where there is no Internet access at all, for instance because of a disaster event. It relies on WiFi to create a ad-hoc peer-to-peer network among devices. Due to this lower level aspect, Serval Mesh requires privileged access to the hardware and some higher skill than Firechat, which is easier to setup but requires the availability of a Internet connection. Serval Mesh has been used [AGS11] to build an alternative, purely peer-to-peer telephony network.

[4] https://opengarden.com/firechat
[5] http://www.whatsapp.com/
[6] https://telegram.org/

Despite the existence of such mesh-oriented applications, however, a general and widespread approach for easily design and program the devices that compose our pervasive continuum is still under investigation [RCKZ13].

2.4 Final remarks

In this chapter we have discussed how quickly technology is evolving, in particular for what concerns communications. We have walked through the recent history of mobile devices, and got a glimpse to the evolving area of wearable devices. We have analysed the most diffused communication technologies: International Mobile Telecommunications, WiFi, Bluetooth and NFC. Finally, we have presented the fascinating idea of considering our world as a pervasive computational continuum, which calls for novel paradigms and engineering procedures and tools. The most consistent contribution of this thesis is devoted to the research of general, well engineered approaches to build such systems.

3
Self-organisation

In this chapter we focus on the software, and in particular on the challenges of engineering the development of software that will run on an ensemble of pervasive devices. We first discuss the issue of coordination and self-organisation: how do we make all those possibly devices collaborate together in order to achieve a global goal, without a centralised decision-maker? Which software platform may we devise to ease this operation? We will see that similar problems have already been successfully solved in nature: the mechanisms underlying such natural behaviours, can, if properly mimicked, help to realise solutions in software systems. We will then run through the existing literature on the issue, analysing the existing platforms supporting pervasive computing, and the tools that can be used to test and debug applications prior to deployment. Finally, we will discuss the shortcomings of the existing technology, and pave the way for the contribution of this PhD thesis.

3.1 Software ecosystems

Regardless the name that we want to use, being it "Internet of Things" rather than "Smart cities" or "Pervasive computing", in all cases we are talking about a system that is expected to host many computations that feature, according to [ZOA+14]:

Situatedness — Pervasive services are typically time- and space-dependent, and feature physically- or socially-situated activities. Components of pervasive systems should be able to interact with the surrounding physical and social world by adapting their behaviour accordingly.

Autonomy and self-adaptivity — While individual components should be autonomous in the face of the inherent dynamics of their operational environment [GPO13], pervasive systems should also feature *system-level autonomy* to deal globally with the unpredictability of the environment, providing properties such as self-adaptation, self-management, and self-organization [MMTZ06].

Prosumption and diversity — Infrastructures for pervasive systems must promote open models of component integration, to be able to take advantage of the injection of new services and components [ZO04]. This is particularly true in the context of *socio-technical systems*

[OC13], where human users and software agents act as *prosumers* – both consumers and producers – of devices, data, and services.

Eternity — As well as short-term adaptation, a pervasive systems infrastructure should allow for the long-term evolution of organizations, components, and patterns of usage, in order to accommodate technological advances as well as the mutable needs of users without requiring extensive re-engineering effort [Jaz05]. In fact, pervasive systems are better conceived as *eternal* systems, engineered for continuous, unlimited service, upgrading, and maintenance over time.

At the same time, "traditional" networks are also increasing in scale and importance for enterprises both large and small. In an increasingly information-dependent and interconnected world, the rising cost of managing and maintaining such systems is driving a search for solutions that can increase the autonomy of computing systems, enabling them to act more as collective services than individual machines [EAWS12, HQL+11], and continuing working over time. In both of these cases, and a number of other areas facing similar challenges (e.g., large-scale sensor networks, multi-UAV control), there is a growing recognition that new paradigms are needed to face the challenge of coordination, namely, engineering the space of interactions [Weg97]. The goal of such engineering paradigms is to find reliable processes that lead to self-organising aggregates that act more as a collective service than as individual machines [CLMZ03, BMP+11, LPF12, SRM+13].

3.2 Taking inspiration from nature

Natural systems are good in dealing with the *complexity* of coordinating large-scale software ecosystems [Sha04, KR08, LT06], since they natively feature key properties such as autonomy, openness, fault tolerance, situated behaviour, self-adaptation and robustness. Many specific models have been proposed in literature, taking inspiration from natural ecosystems, social insects, biochemistry, chemistry and also physics [ZV11].

3.2.1 Physical inspiration

To the category of physical inspired systems belong all those works that take inspiration on the way physical particles move and self-organise according to gravitational and electromagnetic forces. For instance, in TOTA [MZ06a], computational "force fields" – generated either by coordinated components or by the coordination middleware – propagate across a spatial environment, leading to distributed data structures that affect the actions and motions of the agents in that environment. A similar approach is proposed in Co-fields [MZL03], exploiting composite computational fields to coordinate the motion of users and robots in an environment. Physical inspiration also has influenced self-assembly, for instance Guo et al. [GJM12] suggest the adoption of virtual force fields in order to control the material's shape.

A number of "bottom-up" methods implement computational fields using only the local view, including the Hood sensor network abstraction [WSBC04], Butera's "paintable computing" hardware model [But02], TOTA [MZ09], the chemical models [VCMZ11], and Meld [ARGL+07]. The MGS language [GMCS05] takes a notable and different approach, using a method like local field computation to define and evolve the shape of the manifolds on which it executes.

Hybrid automata [Hen96] are particularly interesting for they are relying on a continuous computational space. They are based on the idea of coupling discrete, automata-like descriptions with differential equations that describe continuous evolutions of numerical values, e.g. to consider interaction with sensors/actuators in the physical world. Continuous Spatial Automata [Mac90] push the idea further, by also assuming a continuous spatial substrate over which computation can occur.

More explicit models of field computations are provided by some sensor nework programming models (e.g., Abstract Regions [WM04] and Regiment [NW04]), as well as a number of parallel computing models, most notably StarLisp [LMMD88] and systolic computing (e.g., [EC89, RL93]), which use parallel shifting of data on a structured network.

3.2.2 (Bio)chemical inspiration

Chemical reactions can be seen as an ensemble of myriads of simple laws that generate and regulate the evolution of extraordinarily complex molecular structures. Those rules coordinate in some way the behaviours of a huge number of components, up to the point where they reductionistically drive the evolution of complex assemblies such as biological organisms and meteorological systems.

Gamma [BLM90] was the first and the most prominent example of a chemically-inspired model. In Gamma, coordination is conceived as the evolution of a space governed by chemical-like rules, globally working as a rewriting system [BFLM01]. In the CHAM (chemical abstract machine) model [Ber92a], states are interpreted as chemical solutions, where floating molecules (representing coordinated entities) interact according to some reaction rules, and where *membranes* constrain the execution of reactions.

In chemical tuple spaces [VCNO10] data, devices, and software agents are uniformly represented in the form of chemical reactants, and system behaviour is expressed by means of full-fledged chemical-like laws that are both time-dependent and stochastic rather than in form of rewriting rules. *Biochemical tuple spaces* [VCMZ11] enhance chemical tuple spaces by shaping coordination through distribution and topology. Biochemical coordination models appear very flexible in enabling the spatial formation of both localised and distributed activity patterns, and have been exploited in many special-purpose pervasive frameworks, including crowd mobility management [WBYI08] and participatory sensing [LMG+09]. The amorphous computing model can be considered an example of biochemical coordination model [AAC+00].

Beside having been conceived as general computational model, (bio)chemically-inspired computing can be an effective starting point to realize schemes of dynamic service composition [FDMSS10] or knowledge aggregation [MO13]. Network protocols for data distribution

and aggregation according to chemical models have also been explored, as in the Fraglets approach [MYT07, MMTL13]. Several proposals exist to support service composition based on chemical models [BP09], including proposals specifically conceived for adaptive pervasive services [HAES09], or to support the adaptive organization of knowledge [MO13].

3.2.3 Stigmergy

In the computer science literature, the term "stigmergy" refers to a set of nature-inspired coordination mechanisms mediated by the environment [TB99, Bon99]. Agents deposit data into a distributed, shared, environment so as to collectively (yet implicitly) build distributed data structures that can help them navigate in such environments. The notion of *stigmergy* was introduced as the fundamental coordination mechanism in termite societies [Gra59].

Nowadays, the most widely-studied example of stigmergic coordination in insect societies is probably that of ant colonies [DS04]. The basic mechanism is based on *pheromones* that are released in the environment by the ants that find food on their way back to the nest, thus building pheromone trails towards food that other ants are then stimulated to follow. The pheromones act as environment markers for specific social activities, driving both the individual and the social behaviour of ants. For instance, digital pheromones [PBS02, Par06] have been fruitfully exploited as the basic mechanism for coordinating the movements of robot swarms and modular robots [MZ05], for helping people find directions in an unknown environment [MZ07], and for efficiently cluster information in networks [CVG09, PVM$^+$10]. In the area of networking, stigmergy has been exploited to realize effective routing mechanisms in dynamic networks [BDT99, DDF$^+$06].

3.3 Spatial patterns

Several recent works ([FMDMSM$^+$12b, GVO07, MMTZ06, DWH07]) have focussed on the issue of making self-organisation mechanisms systematically applicable, by proposing descriptions of those mechanisms under the form of catalogues of software design patterns [GHJV95]. In this work, from now on, we will use the word "pattern" to refer to such self-organisation mechanisms, despite the fact that their genesis and applicability is different from the work of Gamma. Recent studies on self-organisation identified existing and new patterns to design computing applications featuring desirable properties of adaptation and robustness to unpredicted changes in the environment, such as human interaction, network mobility and faults, events in the physical environment. As a reference, we consider the work of Fernandez-Marquez et al. [FMDMSM$^+$12b], in which a catalogue of patterns for self-organisation is introduced. Following a series of work in self-organising coordination [VCMZ11, VPMS12], such a catalogue uses a paradigm of chemical-like transformation rules (manipulating information items residing in the network nodes) to specify pattern behaviour and pattern combination. This catalogue is organised in layers (see Figure 3.1): low-level patterns achieve basic self-organisation mechanisms

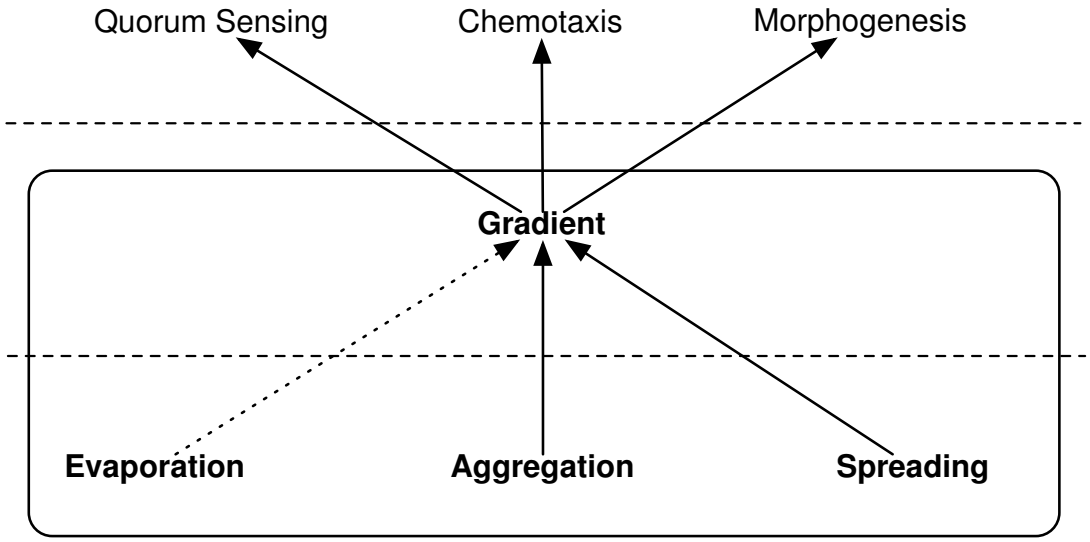

Figure 3.1: Self-organising patterns and their relationships. More fundamental, simpler mechanisms are in the bottom layer, higher level patterns are at the top. The Gradient pattern plays a fundamental role, being the pillar sustaining the higher level patterns. Image adapted from the work of Fernandez-Marquez et al.[FMDMSM12a].

(Spreading, Aggregation and Evaporation), middle-level patterns are used to spatially organise information (gradient), and high-level patterns exploit such information to achieve individual or social goals in a multi-agent system (quorum sensing, chemotaxis, morphogenesis). In the following, we will analyse the gradient pattern. It has a central role, and will be used extensively in the next chapters. We will also quickly discuss Chemotaxis, as an example of upper-level pattern leveraging gradient.

3.3.1 Gradient

A simple though paramount data structure that is a building block of many of the more advanced patterns is the spatial gradient, which assigns to each node a value Γ depending on its position in time and space and on its context [MZ09, BBVT08, VCMZ11]. This structure originates in one or more devices called sources. In every source device, $\Gamma = 0$. In every other device, let N be the set of devices connected with it, and let n be the n-th neighbouring device. For this device, $\Gamma = [\min(f(n))|n \in N]$: namely, the value of the device is the minimum of a function which operates on the neighbours. For instance, if $\forall n\colon\ f(n) = \Gamma_n + 1$, where Γ_n is the value of the gradient in n, then the local value will reflect the minimum hop count towards the nearest source. If $f(n) = \Gamma_n + d(n)$ where $d(n)$ measures the actual distance from the device towards n, then the value of the gradient will approximate the distance from the nearest source. Finally, along with

the local value, a gradient can carry more information, e.g. some strings, the position of n, or numeric values. In this case, the gradient is a "gradcast", namely a mean to transmit an keep an information up-to date in both space and time. This can be understood as a computational field [BDU+13, MZ09, VCMZ11], namely, a distributed data structure mapping each node of the network (or of a subpart of the network) to some (possibly structured) value.

Interestingly, the Gradient pattern can be designed to automatically self-heal to changes in the topology and to movement of the source [Bea09]. An interesting issue with such gradients concerns the asymmetry between how quickly the distributed data structure reacts to changes if they are rising or decreasing the value [BBVT08]. Since the proposed implementation chooses the minimum gradient value from neighbours, it is very quick in detecting and self-healing in case the gradient value should decrease: as soon as a new neighbour discovers a better route towards source, the node updates and gets the correct value. In case the optimal path is broken, instead, at each "computational round" the value rises depending on the distance between a node and the current optimal neighbour. This problem, known as "rising values problem", gets worse with the density of the network[1]. A solution for this problem is available in literature [BBVT08]. In this thesis we will often use the most classic implementation of gradient, and consequently we will run across the rising value problem. The effect of this problem, and its relationship with devices density, are very evident in the experiments of Section 12.4. The curious reader may peek that part to realise of how impacting the effect is.

3.3.2 Chemotaxis

As a main example of upper-level pattern exploiting gradient, chemotaxis is about making a mobile agent (or some information) descending a gradient data structure so as to reach the source from any node the Gradient structure has reached. This is achieved by making the agent continuously reading the estimated distance value as reported in the gradient annotations stored in its neighbourhood, and then moving towards the node exposing minimum distance – the source is eventually reached by construction. This pattern finds applications in long distance communications, or in physical items retrieval in situated applications such as those of pervasive computing.

3.4 Tuple-based coordination

In a tuple-based coordination model [RCD01], software agents synchronize, co-operate, and compete based on *tuples*, which are simple data structures representing information chunks. Tuples are made available in *tuple spaces*, which are shared information spaces working as the *coordination medium*. Coordination occurs by accessing, consuming, and producing tuples

[1] If the density is higher and the number of neighbours is constant, then necessarily there is a smaller distance between nodes. Since the gradient uses distance between nodes to compute its value, this is detrimental to performance.

in an *associative* way, relying on the actual content of tuples and not on any form of naming, addressing, or indexing. An interesting survey of the technologies and platforms [MT03] analyses the suitability of tuple-based coordination systems in supporting openness, unpredictable changes in distributed environments, and several aspects related to adaptiveness: requirements that are of primary importance in pervasive systems.

Several implementations from both academia and industry exist.

Anthill [BMM02] is a framework meant to support design and development of adaptive peer-to-peer applications, in which each node is provided with a local tuple space, agents can travel the network and interact indirectly reading, writing and retrieving tuples.

Biochemical tuple spaces [VC09] is a framework in which a tuple resembles a chemical substance, a notion of activity/pertinency value for tuples is used to model chemical concentration, coordination rules are structured as chemical reactions evolving tuple concentration over time, a tuple space resembles a single-compartment solution, and finally a network of tuple spaces resembles a tissue-like biological system.

CArtAgO [RPV11] is a general purpose framework/infrastructure that makes it possible to program and execute virtual environments - also said virtual / application / software environments - for multi-agent systems.

GigaSpaces XAP [2] (pronounced zap) is a scale-out application server.

JavaSpaces [FHA99, NZ01] are best considered as distributed shared memory with additional features which provide transactional integrity and support for the handling of failure. They are best suited to problems which can be modelled as flows of objects.

Lime [MPR06] is a coordination model and middleware based on the ideas found in Linda, a shared tuple space model developed at Yale in the 1980s. Lime has three primary incarnations, each tailored for a specific environment, ranging from mobile ad hoc networks to sensor networks with base stations or with actuators.

Linda [Gel85] is the common ancestor of every tuple based coordination framework.

MARS [CLZ00] implements a coordination architecture based on Linda-like tuple spaces, associated to each execution environment, where agents can store and retrieve messages in a spatial and temporal uncoupled way. In addition, MARS is programmable, i.e., agents can associate reactions to the operations made on the space.

Molecules of knowledge [MO13] (MoK for short) is a model for knowledge self-organisation, exploiting the biochemical metaphor for its basic abstractions, and biochemical coordination as its coordination model. As far as the basic MoK abstractions are concerned, in MoK

[2] `http://www.gigaspaces.com`

knowledge atoms are generated by knowledge sources in shared spaces – compartments – self-aggregate to shape knowledge molecules, and autonomously move toward knowledge consumers, whose actions (either epistemic or not) are represented as enzymes. As far as the MoK computational model is concerned, MoK features biochemical tuple spaces for the creation, aggregation, diffusion and consumption of knowledge atoms and molecules.

SwarmLinda [TM03] is a middleware that exploits the idea of the collective intelligence, displayed by swarms of ants, for guiding agents in charge of tuple storage and efficient tuple retrieval. Tuples are handled as sort of pheromones or items that ants (agents) relocate in order to improve overall efficiency.

T Spaces [WMLF98] is a tuplespace-based network communication buffer with database capabilities that enables communication between applications and devices in a network of heterogeneous computers and operating systems.

TOTA [MZ09, MZ06b] is a tuple-based middleware explicitly conceived to support field-based coordination for adaptive context-aware and spatially-aware activities in pervasive computing scenarios. In TOTA, each tuple also carries with it a diffusion rule and a maintenance rule in addition to a content. Tota inspired other works, such as the evolving tuples model [SJ07], which adds to the tuples a form of context awareness, in form of possible evolution and adaptation to environmental changes.

TuCSoN [OZ99] is a Java library providing coordination as a service to Java and tuProlog agents. TuCSoN is a model for the coordination of distributed processes, as well as autonomous, intelligent and mobile agents. TuCSoN exploits tuple centres as its coordination media, which are tuple spaces enhanced with the notion of behaviour specification, thanks to the ReSpecT language.

3.4.1 SAPERE

SAPERE[3] [ZCF+11] is a nature-inspired coordination model and framework to support the design and development of composite pervasive service systems. Its reference architecture and coordination model synthesize from existing nature-inspired approaches [VC09, VCO09, VCMZ11, PVM+10], and is based on an assumption of spatial, local interactions (to be realized via a network of distributed tuple spaces), which is in line with all nature-inspired approaches. Its coordination laws make it possible to express and deploy general nature-inspired distributed algorithms and coordination patterns. SAPERE abstracts a pervasive environment as a non-layered *spatial substrate* deployed upon a dense network of connected heterogeneous ICT devices (Figure 3.2). SAPERE acts as a shared coordination medium embodying the basic laws of coordination. Its core components are [VPMS12]:

[3]The model was developed as part of an EU-funded research project. See `http://www.sapere-project.eu/`.

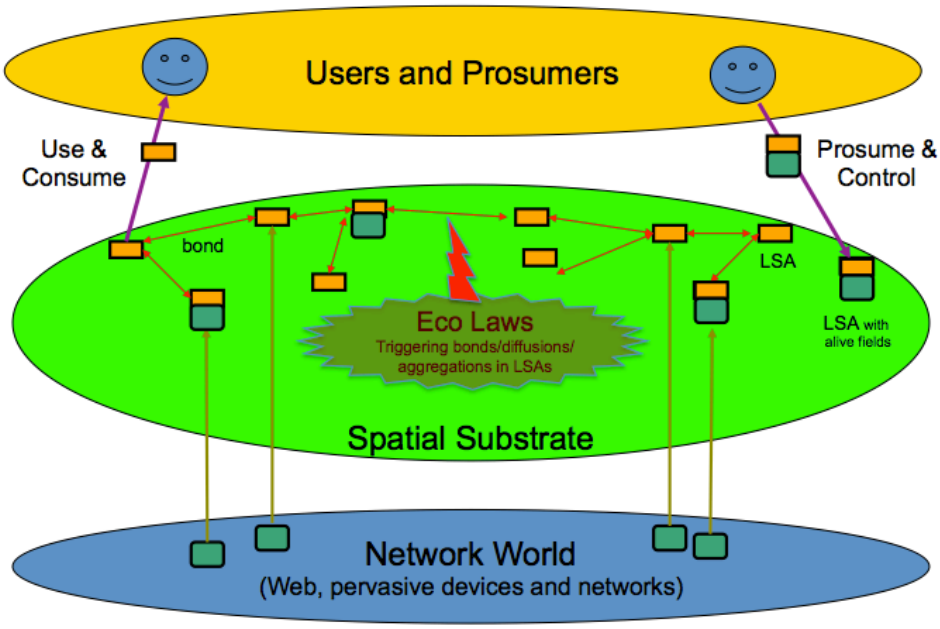

Figure 3.2: The SAPERE reference architecture.

LSA Because of the need to tolerate diversity, a cornerstone of pervasive ecosystems is that a uniform representation is required for the various software agents living within them (whether they run on smartphones, sensors, actuators, displays, or any other computational device). Such a representation needs to expose any information about the agent (state, interface, goal, knowledge) that is pertinent for the ecosystem as a whole or for any subpart of it. In SAPERE this description is called "Live Semantic Annotation" for it should continuously represent the state of its associated component (live), and it should be implicitly or explicitly connected to the context in which such information is produced, interpreted and manipulated (semantic)—possibly relying on standard technologies and techniques of the Semantic Web, like RDF [MM04].

LSA-space To handle situatedness, the behaviour of each agent strictly depends on the local context in which it runs, that is, on the state of other agents living in the same locality (intended as network neighbourhood). As such, the LSAs of each agent are reified in a distributed space (called an "LSA-space") acting as the fabric of the ecosystem, where "context" is simply defined and represented as the set of LSAs stored in a given locality.

LSA bonding Additionally, and in order to make any agent act in a meaningful way with respect to the context in which it is situated, special mechanisms are needed to provide a fine-tuned control of what to each agent is visible/modifiable and what is not. SAPERE tackles this issue by allowing an LSA to include bonds (i.e., references) to other LSAs in the same

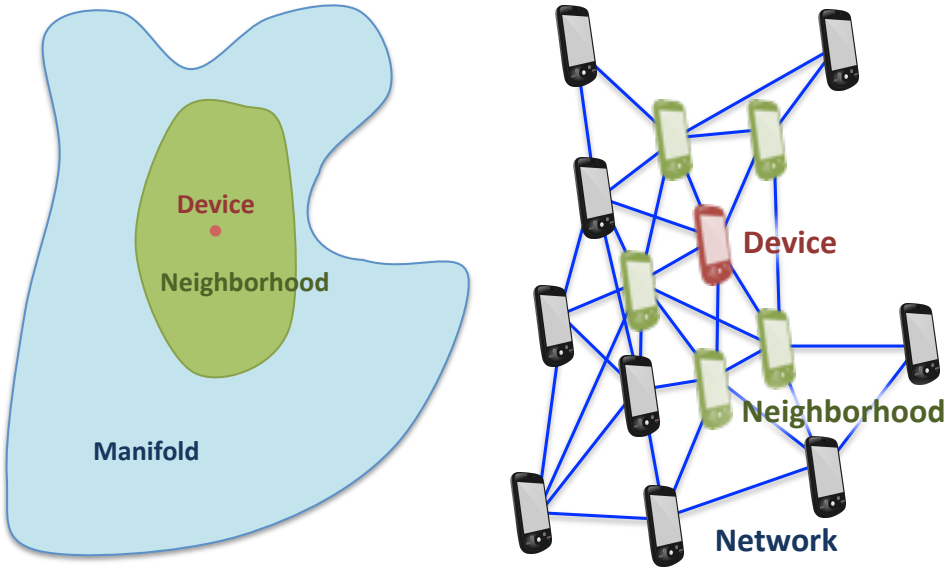

(a) Continuous Space (b) Discrete Network

Figure 3.3: Computational field models originate from approximation of continuous space (a) with discrete networks of devices (b).

context. It is only via a bond that an agent can inspect the state/interface of another agent and act accordingly, while modifications are allowed only to the LSAs an agent injected itself.

Eco-laws Because of adaptivity, while agents enact their individual behaviour by observing their context and updating their LSAs, global behaviour (i.e., global system coordination) is enacted by self-organising manipulation rules of the LSA-space, called eco-laws. They can execute deletion/update/movement/re-bonding actions applied to a small set of LSAs in the same locality. SAPERE structures such eco-laws as chemical-resembling reactions over LSAs—similarly to other approaches like [BP09, VC09, VCMZ11].

3.5 Aggregate programming

Aggregate programming is founded on the observation that in many cases the users of a system are much less concerned with individual devices than with the services provided by the collection of devices as a whole. Typical device-centric programming languages, however, force a programmer to focus on individual devices and their interactions. As a consequence, several different aspects of a distributed system typically end up entangled together: effectiveness and reliability of communications, coordination in face of changes and failures, and composition of behaviours across different devices and regions. This makes it very difficult to effectively design, debug,

maintain, and compose complex distributed applications.

Aggregate programming generally attempts to address this problem by providing composable abstractions separating these aspects:

1. device-to-device communication is typically made entirely implicit, with higher-level abstractions for controlling efficiency/robustness trade-offs;

2. distributed coordination methods are encapsulated as aggregate-level operations (e.g., measuring distance from a region, spreading a value by gossip, sampling a collection of sensors at a certain resolution in space and time); and

3. the overall system is specified by composing aggregate-level operations, and this specification is then transformed into a complete distributed implementation by a suitable mapping.

In short, the key idea behind aggregate programming is to provide languages and APIs that allow a distributed collection of devices to be programmed in terms of their collective behaviours, keeping outside the sight of the software designer details on how such coordination is actually implemented.

A large number of very diverse aggregate programming approaches have been proposed, including abstract graph processing [GGG05], declarative logic [ARGL+07], map-reduce [DG08], streaming databases [MFHH05], and knowledge-based ensembles [NLPT14]—for a detailed review, see [BDU+13].

Despite the fact that all of them are based on viewing the collection of devices as an approximation of continuous space (as depicted in Figure 3.3), most aggregate programming approaches, however, have been too specialized for particular assumptions or applications to be able to address the complex challenges of these emerging environments.

3.5.1 Proto

Proto is a language and middleware [BB06] that exploits field-based coordination to orchestrate the activities of sensor-actuator networks in a aggregate way. Proto is based on the notion of a *computational field*—a map from devices comprising the system to (possibly structured) values, which is treated as unifying first-class abstraction to model system evolution and environment dynamics.

The programming constructs are combined together to form programs, whose semantics is defined in terms of a sequence of synchronous rounds of evaluation by a discrete network of devices called “computational rounds”. In practice, however, there is no requirement for synchrony, and each device can evaluate its own computational rounds independently.

Despite its qualities, Proto still lacks many features expected in a modern programming language and has an implementation encumbered by a number of obsolete considerations that make it difficult to maintain and extend. Moreover, Proto has a complex operational semantics,

that makes it hard to develop mathematical proofs about Proto programs. Such complexity is one of the pillars sustaining the efforts behind Field Calculus (see Section 3.5.2).

3.5.2 Field Calculus

Field Calculus is an attempt to find a unifying model for programming *computational fields* as a generalization of a wide range of existing approaches. It is a fragment of Proto, (see Section 3.5.1), but it also takes inspiration from various other efforts in the literature [MZ09, NW04, VCMZ11, MZ06b, Nag01, Yam07, NW04]. Formalized as the computational Field Calculus [VDB13], this universal language appears to provide a theoretical foundation on which effective general aggregate programming platforms can be built.

Critically, although originally derived from continuous-space concepts, the calculus does not depend on them and is applicable to any network. Field Calculus is expressive enough to be a universal computing model [BVD14] but terse enough to enable a provable mapping from aggregate specifications to equivalent local implementations.

Field Calculus [VDB13] hence provides a key theoretical and methodological foundation for aggregate programming. Its aim is to provide a universal model that is suitable for mathematical proofs of general properties about aggregate programming and the aggregate/local relationship, just as λ-calculus [Chu32] provides for functional programming, π-calculus for parallel programming [Mil99], or Featherweight Java [IPW01] for Java-like object-oriented programming.

In the Field Calculus, everything is a field: computational fields are used to model every aspect of distributed computation, including input from sensors, network structure, environment interactions, distributed computations (e.g. progressive aggregation and spreading processes), and output for actuators. In particular, Field Calculus is constructed using five basic constructs:

1. function definition and evaluation;
2. "built-in" operations for stateless local computation, e.g., addition, multiplication, reading a sensor;
3. a time-evolution construct which allows for stateful computation;
4. a neighbour-value construct that creates a field of values from a device's neighbours;
5. a restriction operator to select which computations to perform in various regions of space and time.

As well as in Proto, in Field Calculus the semantics of programs created by combining such constructs is defined in terms of a sequence of (possibly synchronous) "computational rounds".

The minimal syntax of Field Calculus has allowed its semantics, including proper coherence of device interactions, to be proven correct and consistent [VDB13]. Additionally, despite its definition in terms of discrete semantics, Field Calculus is also space-time universal [BVD14],

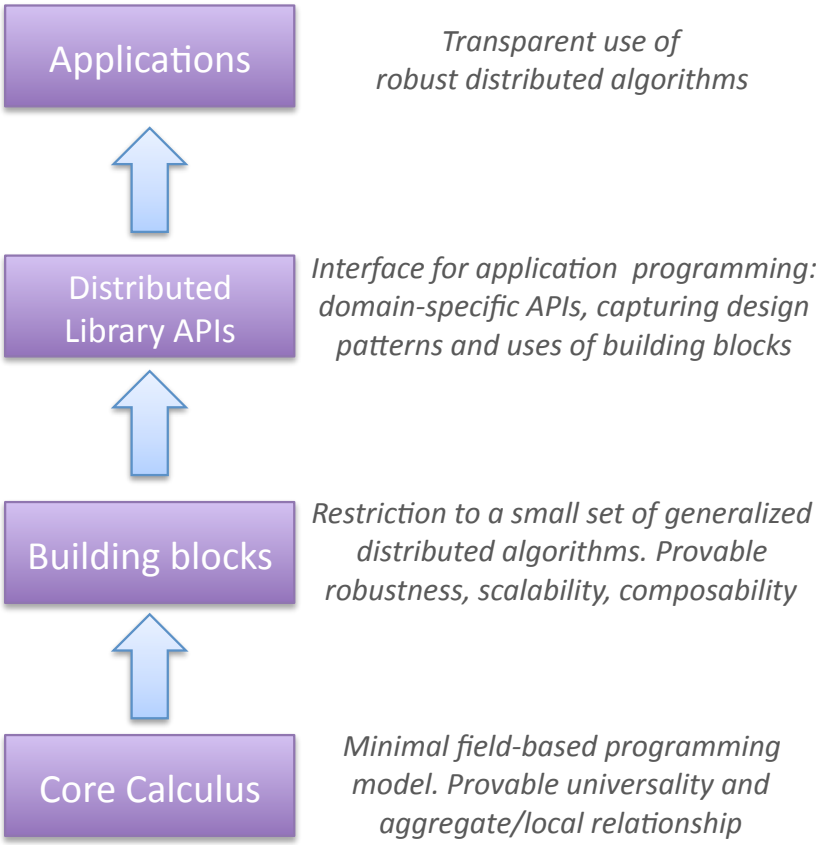

Figure 3.4: Layered approach for development of spatially-distributed systems via aggregate programming.

meaning that it can approximate any field computation, either discrete or continuous, with arbitrary precision given a dense enough network of devices.

This, then, is the key contribution of Field Calculus: any coordination method with a coherent aggregate-level interpretation is guaranteed to be expressible in Field Calculus. Such a method can then be abstracted into a new aggregate-level operation, which can be composed with any other aggregate operation using the rules of built-in functions over fields. Moreover, it can have its space-time extent modulated and controlled by restriction, all while guaranteed that the relationship between global specification and local implementation will always be maintained.

Such a core calculus, like any other, is more a theoretical framework than a practical programming language. In practice, in fact, effective aggregate programming for real-world distributed applications is likely to require a layered approach such as the one depicted in Figure 3.4, of which Field Calculus represents the first block. This kind of layering has been proposed also for works that tackle the same issues with an approach different from aggregate programming, e.g. within the SAPERE project [DMSFMDA14, DAFMDMS14]. Some of the prior approaches on which Field Calculus is based provide very similar semantics (most notably

e	::=	x | l | (b $\overline{\text{e}}$) | (f $\overline{\text{e}}$) | (rep x w e) | (nbr e) | (if e e e)	expression
w	::=	x | l	variable or value
F	::=	(def f($\overline{\text{x}}$) e)	function
P	::=	$\overline{\text{F}}$ e	program

Figure 3.5: Surface syntax of Field Calculus, from [VDB13]

Proto), but they all suffer from some combination of design and implementation problems that render them impractical for widespread adoption. Besides Proto, which is probably the closest to a practical programming environment and whose problems have already been described in Section 3.5.1, others [VPB12], have only minimal implementation.

Basics and Syntax

The basic mechanisms behind Field Calculus can be studied through its syntax, reported in Figure 3.5. We take the global, aggregate-level viewpoint, considering the main syntactic element e as being a field expression, or simply a field. As a standard syntactic notation in calculi for object-oriented and functional languages [IPW01], the overbar notation denotes metavariables over lists, e.g., $\overline{\text{e}}$ ranges over lists of expressions, written $\text{e}_1\ \text{e}_2\ \ldots\ \text{e}_n$.

A basic expression can be a literal value l (also called local value), such as a floating point number, a boolean, or a tuple—note most of the ideas of computational fields are agnostic to the structure of such values. According to the global viewpoint, a literal field expression l actually represents the constant function mapping l to all nodes. A basic expression can also be a variable x, which can be the formal parameter of a function or a store of information to support stateful computations (see rep construct below).

Such basic expressions (values and variables) can be composed by the following 5 constructs.

The first one is *functional composition*, a natural means of manipulating fields as they are functions themselves: $(\text{b}\ \text{e}_1 \text{e}_2\ \ldots\ \text{e}_n)$ is the field obtained by composing together all the fields e_1, e_2, ..., e_n by an operator b. Operators are built-in, and include standard mathematical ones (e.g. addition, sine): they are applied in a pointwise manner to all devices. For instance, if e_t is a field of Fahrenheit temperatures, then the corresponding field of Celsius temperatures is naturally written (* (/ 5 9) (- e_t 32)). Execution of built-in operators is context-dependent, i.e., it can be affected by the current state of the external world. So, 0-ary operator self gives a field that maps each device to its identifier, dt maps each device to the time elapsed since its previous computation round, and finally nbr-range maps each device to a table associating estimated distances to each neighbour (such a table being a field itself).

The second construct is *function (definition and) call*, which is used as abstraction tool and to support recursion: $(\text{f}\ \text{e}_1\ \text{e}_2\ \ldots\ \text{e}_n)$ is the field obtained as result of applying user-defined function f to the fields e_1, e_2, ... e_n. Such functions are declared with syntax (def f($\overline{\text{x}}$) e). For instance, after definition (def convert (x) (* (/ 5 9) (- x 32))), expression

`(convert` e_t`)` denotes the same field of Celsius temperatures as above. Note that function definitions, along with the top-level expression, form a program `P`.

The third construct is *time evolution*, used to keep track of a changing state over time: `(rep x w e)` is initially the field `w` (a local value or a variable) that is stored in the new variable `x`, and at each step in time is updated to a new field as computed by `e`, based on the prior value of `x`. For instance, `(rep x 0 (+ x 1))` is the (evolving) field counting in each device how many rounds that device has computed. Similarly, `(rep x 0 (+ x (dt)))` is the field of time passing.

The fourth construct is *neighbourhood field construction*, the mechanism by which information moves between devices: `(nbr e)` maps each device to the field of its neighbours' local value of field `e`; hence, it is a field of neighbourhood fields like the output of `nbr-range` above. As an example, let `min-hood` be the operator taking a neighbourhood field and returning its minimum value, then `(min-hood (nbr` e_t`))` is the field mapping each device to the minimum temperature perceived in its neighbourhood.

The last construct is *domain restriction*, a sort of distributed branch: `(if` e_0 e_1 e_2`)` is the field obtained by superimposing field e_1 computed everywhere e_0 is true and e_2 everywhere e_0 is false. As an example `(if` e_{fah} e_t `(convert` e_t`))` is the field of temperatures provided in Fahrenheit (resp. Celsius) where the field e_{fah} is true (resp. false). Restriction is the most subtle of the five mechanisms, because it has the effect of preventing the unexpected spreading of computation to devices outside of the required domain, even within arbitrarily nested function calls, as will be clarified in the following.

A simple example illustrates how these five key mechanisms can be combined to implement useful spatial patterns.

```
(def gossip-min (source) (rep d source (min-hood (nbr d))))

(def distance-to (source)
  (rep d infinity (mux source
     0
     (min-hood (+ (nbr d) (nbr-range)))
     )))

(def distance-obs-to (source obstacle)
      (if (not obstacle) (distance-to source) infinity))
```

Function `gossip-min` takes a `source` field and produces a new field mapping each device to the minimum value that `source` initially takes. The `rep` construct initially sets the output variable `d` at `source`, and it iteratively updates the value at each device with the minimum one available in `d`'s neighbours. Hence, `gossip-min` describes a process of gossiping values until the minimum one converges throughout the network.

Similarly, function `distance-to` takes as its input a `source` field holding boolean values, and returns a new scalar field that maps each device to the estimated distance to the nearest

device where `source` is true. This works by first setting `d` to infinity, then updating it as follows: sources are of course at distance 0, while all other devices use the triangle inequality, finding the minimum sum of a neighbour's estimated distance `d` and the distance to that neighbour. Operator `mux`, used to combine the two, is a purely functional multiplexer, which uses the first input to choose whether to return the second or third. The field returned by `distance-to` is a *gradient*.

The last definition exemplifies the use of construct `if`. It creates two different spatial domains: one where the obstacle is present (field `obstacle` holds positive boolean value) and one where is not. In the former an infinity constant field is computed; in the latter we spread the `distance-to` field. As a result, distance estimation as provided by `distance-to` automatically takes into account the need of circumventing obstacle areas, since information does not cross the two domains due to the semantics of `nbr`.

Semantics

The computational Field Calculus formalisation is set forth in Figure 3.6. For a complete description of the symbols and their meaning, the interested reader should refer to [VDB13].

The most important point for the discussion of this thesis is the compactness of such core semantics: it can be summarised in a single page of rules, while several pages would be required to host the complete semantics of Proto [VBC11, VBU13], the language from which Field Calculus is derived. The number of pages required to host a formal semantics is clearly a rule of thumb, but gives also the less experienced reader some idea on how intricate is Proto with respect to Field Calculus. A more lightweight semantics has two key advantages:

1. many properties become mathematically treatable and can be and proved;

2. provides a solid guideline for practical languages to be designed and implemented.

The first point is very relevant for this thesis, since a solid mathematical background is the best possible foundation for a reliable engineering procedure. The second point will be exploited in Chapter 22, where a new practical language building on Field Calculus will be presented. Such language's interpreter adheres to the operational semantics defined in Figure 3.6, and consequently shares all the nice properties of Field Calculus [VDB13].

3.6 Engineering and tools

Despite the efforts of addressing specific aspects of self-adaptability and openness in pervasive contexts [MMTZ06, BCD+06, KR08], yet a comprehensive engineering framework is missing. The hardest challenge is due to the fact that most of the complexity of such systems does not come from the individual behaviours, but rather from interaction [Weg97, GSW06, ORV06].

Runtime Expression Syntax:

$$
\begin{array}{lcll}
e & ::= & a{\cdot}\mathring{v} & \text{runtime expression (rte)} \\
a & ::= & \mathtt{x} \mid v \mid (\mathtt{nbr}\ e) \mid (\mathtt{if}\ e\,e\,e) \mid (\mathtt{rep}^{\mathtt{s}}\ \mathtt{x}\ \mathtt{w}\ e) \mid (\mathtt{f}^{\mathtt{s}}\ \overline{e}) \mid (\mathtt{b}\ \overline{e}) & \text{auxiliary rte} \\
v & ::= & \mathtt{l} \mid \phi & \text{runtime value} \\
\mathtt{s} & ::= & \mathring{a} & \text{superscript} \\
\mathtt{w} & ::= & \mathtt{x} \mid \mathtt{l} & \text{variable or local value} \\
\phi & ::= & \overline{\sigma} \mapsto \overline{\mathtt{l}} & \text{field value} \\
\Theta & ::= & \overline{\sigma} \mapsto \overline{e} & \text{tree environment} \\
\Gamma & ::= & \overline{\mathtt{x}} := \overline{v} & \text{variable environment}
\end{array}
$$

Congruence Contexts:

$$
\mathbb{C} ::= (\mathtt{nbr}\ [\,]) \mid (\mathtt{f}^{\mathtt{s}}\ \overline{e}\ [\,]\ \overline{e}) \mid (\mathtt{b}\ \overline{e}\ [\,]\ \overline{e}) \mid (\mathtt{if}\ [\,]\ e\ e) \mid (\mathtt{if}\ a{\cdot}t\ [\,]\ e) \mid (\mathtt{if}\ a{\cdot}f\ e\ [\,])
$$

Alignment contexts:

$$
\mathbb{A} ::= \mathbb{C} \mid (\mathtt{rep}^{\mathtt{s}}\ \mathtt{x}\ \mathtt{w}\ [\,]) \mid (\mathtt{f}^{[]}\ \overline{a{\cdot}v})
$$

Auxiliary functions:

$$
\begin{array}{rcll}
\pi_{\mathbb{A}}(\Theta,\Theta') & = & \pi_{\mathbb{A}}(\Theta),\pi_{\mathbb{A}}(\Theta') & \\
\pi_{\mathbb{A}}(\sigma \mapsto (\mathbb{A}'[e]){\cdot}v) & = & \sigma \mapsto e & \text{if } \mathbb{A}' :: \mathbb{A} \\
\pi_{\mathbb{A}}(\sigma \mapsto e) & = & \bullet & \text{otherwise}
\end{array}
$$

$$
\mathtt{s} \rhd a = a \qquad \mathtt{s} \rhd \circ = \mathtt{s}
$$

$$
\begin{array}{rcl}
(\mathtt{nbr}\ [\,]) & :: & (\mathtt{nbr}\ [\,]) \\
(\mathtt{f}^{\mathtt{s}'}\ e'_1 ... e'_{i-1}\ [\,]\ e'_{i+1} ... e'_n) & :: & (\mathtt{f}^{\mathtt{s}}\ e_1 ... e_{i-1}\ [\,]\ e_{i+1} ... e_n) \\
(\mathtt{b}\ e'_1 ... e'_{i-1}\ [\,]\ e'_{i+1} ... e'_n) & :: & (\mathtt{b}\ e_1 ... e_{i-1}\ [\,]\ e_{i+1} ... e_n) \\
(\mathtt{if}\ [\,]\ e'_1\ e'_2)) & :: & (\mathtt{if}\ [\,]\ e_1\ e_2)) \\
(\mathtt{if}\ a'{\cdot}t\ [\,]\ e') & :: & (\mathtt{if}\ a{\cdot}t\ [\,]\ e) \\
(\mathtt{if}\ a'{\cdot}f\ e'\ [\,]) & :: & (\mathtt{if}\ a{\cdot}f\ e\ [\,]) \\
(\mathtt{rep}^{\mathtt{s}'}\ \mathtt{x}\ \mathtt{w}\ [\,]) & :: & (\mathtt{rep}^{\mathtt{s}}\ \mathtt{x}\ \mathtt{w}\ [\,]) \\
(\mathtt{f}^{[]}\ e'_1 ... e'_n) & :: & (\mathtt{f}^{[]}\ e_1 ... e_n)
\end{array}
$$

Reduction Rules:

$$
\frac{}{\Theta;\Gamma \vdash v \to v{\cdot}v}\ [\text{VAL}]
$$

$$
\frac{}{\Theta;\Gamma \vdash \mathtt{x} \to \mathtt{x}{\cdot}\Gamma(\mathtt{x})|_{dom(\Theta),\varepsilon(\mathtt{self})}}\ [\text{VAR}]
$$

$$
\frac{\pi_{(\mathtt{nbr}\ [])}(\Theta) = \overline{\sigma} \mapsto \overline{a{\cdot}\mathtt{l}} \qquad \phi = (\overline{\sigma} \mapsto \overline{\mathtt{l}}, \varepsilon(\mathtt{self}) \mapsto \mathtt{l})}{\Theta;\Gamma \vdash (\mathtt{nbr}\ a{\cdot}\mathtt{l}) \to (\mathtt{nbr}\ a{\cdot}\mathtt{l}){\cdot}\phi}\ [\text{NBR}]
$$

$$
\frac{}{\Theta;\Gamma \vdash (\mathtt{b}\ \overline{a{\cdot}v}) \to (\mathtt{b}\ \overline{a{\cdot}v}){\cdot}\varepsilon(\mathtt{b},\overline{v})}\ [\text{OP}]
$$

$$
\frac{}{\Theta;\Gamma \vdash (\mathtt{if}\ a{\cdot}t\ a'{\cdot}\mathtt{l}\ e) \to (\mathtt{if}\ a{\cdot}t\ a'{\cdot}\mathtt{l}\ |e|){\cdot}\mathtt{l}}\ [\text{THEN}]
$$

$$
\frac{}{\Theta;\Gamma \vdash (\mathtt{if}\ a{\cdot}f\ e\ a'{\cdot}\mathtt{l}) \to (\mathtt{if}\ a{\cdot}f\ |e|\ a'{\cdot}\mathtt{l}){\cdot}\mathtt{l}}\ [\text{ELSE}]
$$

$$
\frac{\pi_{\mathbb{C}}(\Theta);\Gamma \vdash a \to e}{\Theta;\Gamma \vdash \mathbb{C}[a] \to \mathbb{C}[e]}\ [\text{CONG}]
$$

$$
\frac{\pi_{(\mathtt{rep}^{\mathring{1}}\ \mathtt{x}\ \mathtt{w}\ [])}(\Theta);\Gamma,(\mathtt{x} := (\Gamma(\mathtt{w}) \rhd \mathring{1})) \vdash a \to a'{\cdot}\mathring{v}}{\Theta;\Gamma \vdash (\mathtt{rep}^{\mathring{1}}\ \mathtt{x}\ \mathtt{w}\ a) \to (\mathtt{rep}^{\mathring{1}\rhd\mathring{v}}\ \mathtt{x}\ \mathtt{w}\ a'{\cdot}\mathring{v}){\cdot}\mathring{v}}\ [\text{REP}]
$$

$$
\frac{\pi_{(\mathtt{f}^{[]}\ \overline{a{\cdot}v})}(\Theta);(args(\mathtt{f}) := \overline{v}) \vdash (body(\mathtt{f}) \rhd \mathtt{s}) \to a{\cdot}\mathring{v}}{\Theta;\Gamma \vdash (\mathtt{f}^{\mathtt{s}}\ \overline{a{\cdot}v}) \to (\mathtt{f}^{a}\ \overline{a{\cdot}v}){\cdot}\mathring{v}}\ [\text{FUN}]
$$

Figure 3.6: Device Operational Semantics

The emergence of such behaviour poses a major issue: if we focus on the single interacting part of the whole system, we encounter a local-to-global problem. No methodology yet successfully solved the task of describing the behaviour of the aggregate and compile it to the program that each device should run, although a very promising research line is the one of aggregate programming [BB06, VDB13]. One of the problems under investigation is the possibility of defining high level constructs based on well known building blocks. This thesis will outline some novel work in this direction in Chapter 21 and Chapter 22.

3.6.1 Simulation

Multiple methodologies have been proposed to tackle complex pervasive systems [MCOV11, BW08, BMV09a, MN10, ZCF+11], and they always include *simulation* as a key step to realise what-if analysis prior to actual development, to both assess the general validity of the designed mechanisms and to fine tune system parameters. There are many kinds of simulation tools available: they either provide programming/specification languages devoted to ease construction of the simulation process, especially targeting computing and social simulation (e.g. as in the case of multi-agent based simulation [LCRP+05, BMV09a, SGJ07, BMV09b, NHCV07, Skl07]), or they stick to quite foundational computing languages to better tackle performance, mostly used in biology-oriented applications [Pri95, Mur89, UP05, EMRU07].

None of the existing tools, however, aims at bridging the gap between these approaches, trying to extend a basic computing model – still retaining high performance – toward ease applicability to complex situated computational systems. In this thesis, a novel toolchain for engineering pervasive systems will be presented in Part II. The idea behind this tool is to rely on the basic model of chemical reactions, whose stochastic simulators are known for their high performance, and try to extend the existing algorithms in literature up to the point that richer models can be supported. The tool also includes a probabilistic model checker, which can be leveraged to estimate with the desired precision the value of model properties that would be otherwise hard to evaluate.

3.7 Final remarks

In this chapter, we have focussed on the state of the art in self-organisation. We have described the main features of a software ecosystem (namely: situatedness, autonomy and self-adaptivity, prosumption and diversity, eternity) seeing how most of the existing works take inspiration from nature (either from physics, chemistry, biochemistry or stigmergy) in order to deliver self-* properties. We have discussed the existing efforts to extract from those approaches a catalogue of reusable mechanisms, named spatial patterns (taking inspiration from the work of Gamma et al. [GHJV95], even though they refer to a different kind of patterns). We described the most important of them, the gradient, which is used by higher level patterns as a mechanism. Two of the most prominent approaches to the organisation of a computational ecosystems

were presented: tuple based coordination (with a deeper analysis of the Self-Aware Pervasive Ecosystems approach) and aggregate programming (with an analysis of Proto and its distillate Field Calculus). Finally, the importance of simulation and the lack of comprehensive tools and procedures for engineering computational ecosystems was discussed.

Part II

An integrated toolchain for pervasive ecosystems

4

Chemical-inspired engine

This part of the thesis is devoted to describing a novel toolchain whose final goal is to provide proper support to engineers of pervasive ecosystems.

As suggested in the introduction to this work, such a software must provide a very generic meta-meta model, that will need to be instanced in a meta-model that exposes the abstractions the designer intends to use in order to build the final system. This meta-model is then instanced into a specific model (e.g., if the meta-model is a programming language, then the model is a program); and must be tested and verified. In order to do so, the model must be supported by a formal description of the environment: being the system evolution strongly influenced by the context in which it is deployed, the environment must be considered in the testing phase.

As previously discussed, simulation is often the only viable solution to the analysis of such complex system. In order to run simulations on the model, a simulation engine is required. Moreover, if we want to reuse such engine for every meta-model, it must be flexible enough to simulate any instance of the meta-meta model.

Once we have a complete simulation solution, a question arises: how many simulations should the designer run in order to measure some property with the desired confidence and approximation? The answer will be given in this part, and an automation solution will be proposed chaining the simulator described above with an existing stochastic model checker.

In the following, we present a discrete event simulation engine derived from the popular and successful Gillespie's stochastic simulation algorithm [Gil77], and in particular from its notable and more efficient extensions developed by Gibson-Bruck [GB00] and by Slepoy [STP08]. Gillespie's SSA is a discrete and stochastic method that intrinsically owns the event-driven properties [BBL$^+$11]: introduced to model chemical systems, nowadays it represents the basis of many simulation platforms, in particular those developed for the investigation of biochemical systems [PRSS01, Kie02, CG09, VB08, MV10, GPOS13, HSG$^+$06]. With significative extensions, it was recently used to model artificial systems grounding on computational models inspired to chemical natural systems and ecology [MVFM$^+$13]. Moreover, its optimised extensions [GB00, STP08] are very efficient, thus allowing really fast simulation runs.

4.1 Gillespie's SSA as an event-driven algorithm

First of all, we summarise the idea that the Gillespie's SSA is grounded on: a chemical system is modelled as a single space filled with molecules that may interact through a number of reactions describing how they combine. The instantaneous speed of a reaction is called propensity and depends on the kinetic rate of the reaction and on the number of molecules of all the reagents involved. For a reaction i of the form:

$$R_0 + R_1 \xrightarrow{r} P_0 + P_1 + \ldots + P_n$$

the propensity a_i is defined as:

$$a_i = r \cdot [R_0] \cdot [R_1]$$

where $[M]$ is the number of molecules of species M. Given that, the algorithm relies on the idea that the system can be simulated by effectively executing the reactions one by one and changing the system status accordingly. Every algorithm follows four main steps:

1. select the next reaction μ to be executed;
2. calculate the time of occurrence of μ according to an exponential time distribution and make it the current simulation time;
3. change the environment status in order to reflect this execution;
4. update the propensities of the reactions.

Such algorithm is event-driven in that it executes only one reaction/event at a time: it changes the state of the system and consequently the event list *i.e.*. which other reactions/events can be executed from there.

4.2 Gillespie's optimised versions

This algorithm has been improved in various works in literature, particularly notable is the work of Gibson and Bruck [GB00] and Slepoy et al. [STP08]. Both works optimise the base algorithm in two phases: the selection of the next reaction to execute and the update of the reaction pool – pending event list – once an event has been executed. For the latter, they both rely on the concept of "dependency graph". A dependency graph (DG) is a statically created directed graph in which nodes are all the reactions in the simulated system, and arcs connect a reaction `r` to all those that depend on it, namely, those whose triggering time should be updated as r is executed. For instance, if `r` changes the concentration of some molecule `m`, all those reactions that use `m` are to be properly re-scheduled as soon as `r` is fired. Even though this optimisation does not affect the execution time in the worst case, it offers great benefits in the average case, since most of the reactions are not interdependent.

The selection of the next reaction to execute is where those improved algorithms differ most. In Slepoy's work, the author divides reactions in groups based on their propensity: if p_{min} is the lowest propensity, the first group contains those reactions whose propensity ranges from p_{min} and $2p_{min}$, the second those between $2p_{min}$ and $4p_{min}$, and so on. Give those groups, an algorithm called "Composition-Rejection" (CR) is applied. In the composition phase, a group is randomly selected in logarithmic time; in the rejection phase a reaction is chosen within a group by throwing two random numbers: the first one, between 0 and the number of reactions in the group, identifies a reaction; the second one, between 0 and the maximum propensity allowed for the group, is used to decide whether or not to actually execute the selected reaction. In case the reaction is rejected, the rejection procedure is re-applied. Given the way groups are built, there is at least a 0.5 probability that the rejection procedure concludes at each attempt. On average, the CR algorithm requires five random numbers to be thrown. This algorithm requires logarithmic time to select the next reaction with respect to the total number of groups, but the author argues which, since in most biological scenarios this number is constant, then the algorithm runs effectively in constant time. In Gibson-Bruck, the same selection operation is made by computing a putative execution time for each reaction, and then using a binary heap data structure to sort them. This way, the next reaction to execute is always the root node of the tree, and the selection is performed in constant time. However, once the reaction has been executed, all the dependent reactions must be updated and re-sorted in the tree. In the worst case, this takes logarithmic time with respect to the number of reactions.

4.3 From SSA to full fledged DES

Both the Gisbon-Bruck's and Slepoy's schedulers are granted to correctly simulate a Poisson process. However, in order to build a full-fledged Discrete Event Simulation (DES) engine, we must offer the possibility to schedule also non-Markovian events. Imagine, for instance, that we want to simulate an agent that does an action every fixed time interval (e.g. a man walking). This, clearly, is not a memoryless process. We argue that Gibson-Bruck offers a more suitable base for a general purpose DES: in fact, its next reaction choosing mechanism is orthogonal to the way the putative times are computed. This intrinsic feature allows to neatly separate the generation of times (namely, the time distribution of each event) and the actual scheduling (choice of which event should run next). We chose the Next Reaction Method, and, consequently, the main simulation algorithm follows the basic steps listed in Algorithm 1.

A second feature that we need in order to shift the paradigm from pure chemistry towards higher expressiveness is the possibility to simulate multiple, separate, interacting and possible mobile entities. This requirement can be partly addressed by the notion of intercommunicating compartments [CG09, VB08, MV10, GPOS13], in a way that allows to also model systems characterised by a set of connected volumes and not only a unique chemical solution. The hardest challenge in simulating multiple compartments with an efficient SSA is in improving the dependency graph: reactions that could be interdependent but happen on separated compartments

Algorithm 1 Simulation flow

```
1: cur_time = 0
2: cur_step = 0
3: for each node n in environment do
4:    for each reaction nr in n do
5:       generate a new putative time for nr
6:       insert nr in DIPQ
7:       generate dependencies for nr
8: while cur_time < max_time and cur_step < max_step do
9:    r = the next reaction to execute
10:   if r's conditions are verified then
11:      execute all the actions of r
12:      for each reaction rd which depends on r do
13:         update the putative execution time
14:   generate a new putative time for r
```

should not be marked as interconnected within the dependency graph. Mobility makes everything even harder, since it may lead to the creation and disruption of communication channels between compartments, and consequently to simulation-time changes in the structure of such dependency graph. Summarising, supporting dynamic environments with multiple mobile compartments require the dependency graph to become a dynamic data structure, which cannot be pre-computed at the simulation initialisation and kept static.

4.3.1 Dynamic Dependency Graph

There are multiple ways to conceive such dynamic data structure. Ideally, it would be possible to just drop the optimisation or reuse the classic definition, in which case the triggering of a reaction in whichever compartment would cause the recalculation of the status of each potentially dependent event in every compartment. This approach would lead to a massive performance impact, since the dependency computation is the most expensive operation. As evidence, some performance analysis [STP08] show the difference between Gibson-Bruck and Slepoy et al. algorithms with and without a dependency graph: the result is very strongly in favour of the former.

A possibility for efficiently adapting a dependency graph to a network of compartments could be to define the input and output contexts for each reaction, namely the places where reactions respectively "read" their reactants and "write" their products. Multiple contexts could be defined, we propose to adopt three levels: `local`, `neighborhood` and `global`. In a purely chemical simulation, all the reactions have a `local` input context and may have either `neighborhood` or `local` output context, depending on whether or not they send molecules towards other compartments.

Let $dep : R^2 \longrightarrow boolean$ be the function that is used to build the static dependency graph in every SSA: given two reactions `r1` and `r2`, $dep(\texttt{r1}, \texttt{r2})$ returns `true` if the propensity of `r2` may be influenced by the execution of `r1`. In a multi-compartment scenario, `r1` influences `r2` (there is an oriented edge in the dependency graph connecting `r1` to `r2`) iff $dep(\texttt{r1}, \texttt{r2})$ and at least one of the following are `true`:

- `r1` and `r2` are on the same node;
- `r1`'s output context is `global`;
- `r2`'s input context is `global`;
- `r1`'s output context is `neighborhood` and `r2`'s node is in `r1`'s node neighbourhood;
- `r2`'s input context is `neighborhood` and `r1`'s node is in `r2`'s node neighbourhood;
- `r1`'s output context and `r2`'s input context are both `neighborhood` and the neighbourhoods of their nodes have at least one common node.

The filters listed above greatly compact the number of edges of a dependency graph in most scenarios, with great benefits on the engine performance. On top of this finer-grain locality concept, if the model supports compartment mobility, the dependency graph must support the dynamic addition and removal of reactions.

Adding a new reaction implies to verify its dependencies against every reaction of the system. In case there is a dependency, it must be added to the dependency graph. Removing a reaction `r` requires to delete all dependencies in which `r` is involved both as influencing and influenced. Moreover, in case of a change of system topology, a dependency check among reactions belonging to nodes with modified neighbourhood is needed. It can be performed by scanning them, calculating the dependencies with the reactions belonging to new neighbours and deleting those with nodes which are no longer in the neighbourhood.

4.3.2 Dynamic Indexed Priority Queue

An issue that arises with addition and removal of nodes from the simulation is the possible unbalancing of the scheduling queue, that in the original work is realised as a binary tree of reactions, whose main property is that each node stores a reaction whose putative time of occurrence is lower than each of its sons.

Our idea is, for each node, to keep track of the number of descendant per branch, having in such way the possibility to keep the tree balanced when adding nodes. In Figure 4.1 we show how the same IPQ drawn in Gibson and Bruck's work [GB00] would appear with our extension. Given this data structure, the procedures to add and remove a new node `n` are described respectively in Algorithm 2 and Algorithm 3, in which the procedure `UPDATE_AUX(n)` is the same described by Gibson and Bruck [GB00].

Using the two procedures described above, the topology of the whole tree is constrained to remain balanced despite the dynamic addition and removal of reactions.

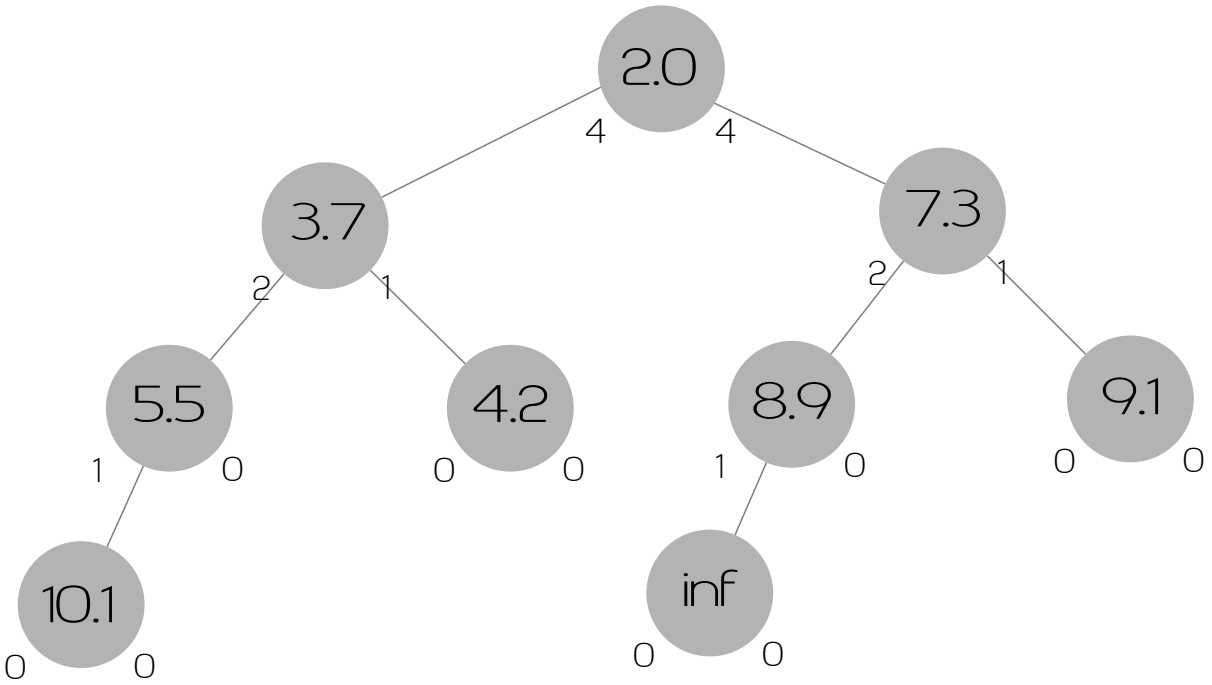

Figure 4.1: Indexed Priority Queue extended with descendant count per branch

Algorithm 2 Procedure to add a new node `n`

```
1:  if root does not exist then
2:      n is the new root
3:  else
4:      c ← root
5:      while c has two descendants do
6:          if c.right < c.left then
7:              dir ← right
8:          else
9:              dir ← left
10:         add 1 to count of dir descendants
11:         c ← c.dir
12:         if c has not the left child then
13:             n becomes left child of c
14:             set count of left nodes of c to 1
15:         else
16:             n becomes right child of c
17:             set count of right nodes of c to 1
18:     UPDATE_AUX(n)
```

Algorithm 3 Procedure to remove a node `n`

```
1:  c ← root
2:  while c is not a leaf do
3:     if c.left > c.right then
4:        dir ← left
5:     else
6:        dir ← right
7:     subtract 1 to count of dir descendants
8:     c ← c.dir
9:  if c ≠ n then
10:    swap n and c
11:    remove n
12:    UPDATE_AUX(c)
13: else
14:    remove n
```

4.4 Final remarks

This chapter introduced a novel the engine for our toolchain targeting the engineering of pervasive ecosystems. We first analysed the existing literature, evaluating pros and cons of the famous extensions to the original Gillespie's stochastic simulation algorithm. We argue that Gibson/Bruck algorithm provides a suitable foundation on which a more general simulation algorithm can be built.

We extended such algorithm in order to be able to simulate more complex scenarios, and in particular to be able to deal with multiple and intercommunicating compartments that may join and leave the system, and whose interconnections may change with time. To the best of our knowledge, no previous attempt was made to extend existing chemical-oriented engines in this direction.

5
Meta-meta model

In this chapter, we will introduce a meta-meta model that the engine described in the previous chapter would be able to simulate. Such meta-meta model is a sort of extended chemistry, that we believe is able to capture the complexity of computational ecosystems. Such complexity is achieved by the following set of common key properties:

- situatedness – they deal with spatially- and possibly socially-situated activities of entities, and should therefore be able to interact with a limited portion of the surrounding world and contextualise their behaviour accordingly;
- adaptivity – they should inherently exhibit properties of autonomous adaptation and management to survive contingencies without external intervention, global supervision, or both;
- self-organisation – spatial and temporal patterns of behaviour should emerge out of local interactions and without a central authority that imposes pre-defined plans.

5.1 Limitations of pure chemistry

Among the many natural metaphors one can use as inspiration for modelling and developing artificial systems with the above properties [ZV11], we consider chemistry following a series of work in the field of pervasive computing [VCMZ11, VZ10, ZCF+11]. We argue that there are three main issues to be resolved in order to build a meta-model that can be sufficiently expressive for our purpose starting from a purely chemical model:

1. the concept of environment where agents are situated and can move is missing in a model that considers only intercommunicating chemical compartments;
2. the only available mean for changing the system status is the execution of a reaction;
3. the only data item that chemical reactions can manipulate are molecules' concentrations, namely numbers connected to a particular token.

5.2 Environment and connection model

In the following discussion, we will use interchangeably compartment/agent/node and reaction/events as synonyms. As first step, we introduce the environment, absent in chemistry-derived SSAs, as first class abstraction. The environment has responsibility to provide, for each compartment, a set of compartments that are its neighbours. The rule which is applied to determine whether or not a node belongs to another node's neighbourhood can be arbitrarily complicated. Also, it is responsible of exposing possible physical boundaries, namely, to limit the possible movements of compartments situated within the environment.

5.3 Reactions, conditions, actions

The fact that reactions are the only abstraction the modeller can rely upon in order to let the simulated system progress is not a difficult problem by itself. In fact, nothing prevents to widen the generality of a reaction by defining it as: "a set of conditions that, when matched, trigger a set of actions on the environment".

With this definition in mind, a condition is a function that associates to each possible state of the environment a numeric value ranging from zero to positive infinity. If such value is zero, the event can not be scheduled; otherwise, it is up to the reaction to interpret the number: it can influence or not the time at which the reaction will be scheduled, depending on the specific reaction implementation. In case we desire to re-build the original chemical model, we would define a condition for each of the molecules on the left-hand side of the chemical reaction that return the number of molecules currently available in the local compartment. Also, we would define the reaction in such a way that it correctly interprets the number returned by the conditions as concentration of each reactant and correctly applies the rules for computing a propensity to be used to influence the reaction speed [Gil77].

In this framework, actions are arbitrary changes of the environment. In case of pure chemistry, the actions of a reaction would be one for each reactant (that must be removed from the local compartment) and one for each product (that must be added to the local compartment). In case of an extended model considering also multiple compartments, an action should be programmed to be responsible of transferring molecules from a node to a neighbouring one.

Both conditions and actions must expose the set of possible data items (molecules) that they may read or modify: this is necessary in order to allow the dependency graph to be built. Also, both conditions and actions must expose a context of the type `local`, `neighborhood` or `global`; it will be used internally to determine the input and output contexts for the reaction itself.

The reactions are responsible of computing their expected execution time. The engine may require such putative time to be updated in two cases: i) the reaction has just been executed or ii) a reaction on which this reaction depends on has been executed. In case of update required, the reaction should leverage a separately defined time distribution to compute the next putative

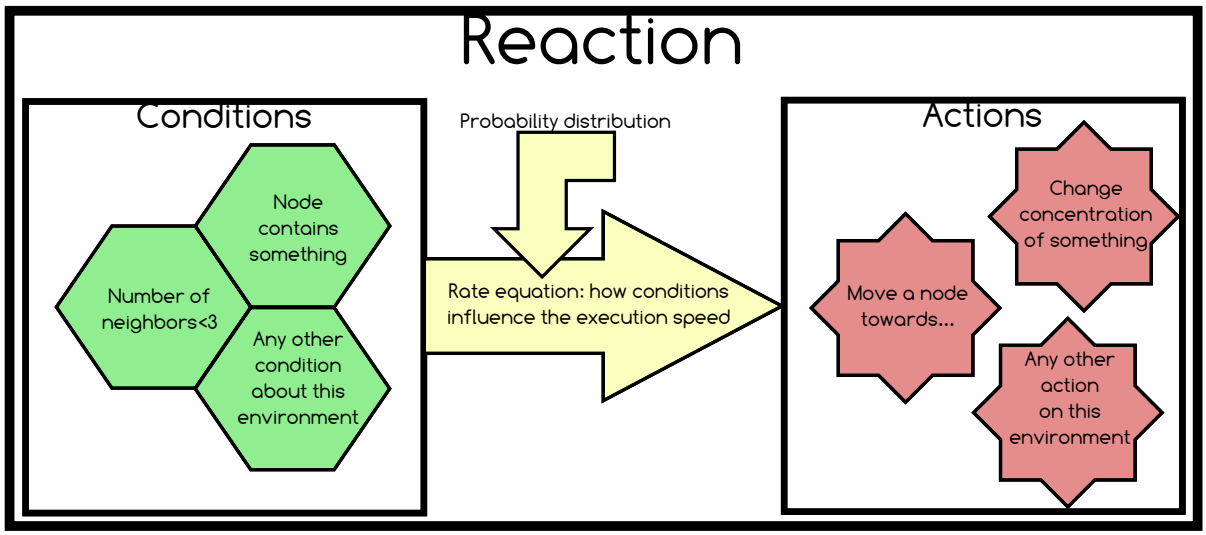

Figure 5.1: Reaction model: a set of conditions on the environment that determines whether the reaction is executable, a rate equation describing how the speed of the reaction changes in response to environment modifications, a probability distribution for the event and a set of actions, which will be the neat effect of the execution of the reaction.

execution time, possibly feeding the time distribution with a summary of the data gathered from conditions. In case of a Poisson process, a negative exponential time distribution initialised with $\lambda = r$ should be used, for instance. In case of a repetitive event, such as a timer, a Dirac comb may be used.

The concept of reaction is graphically depicted in Figure 5.1. It allows, for example, to model reactions which are faster if a node has many neighbours, or also reactions that resemble complex biological phenomena such as the diffusion of morphogenes during embryo development as described in the work of Lecca et al. [LIDP10].

In this model the atomicity of the reactions represent a double edged sword: on the one hand, they allow for arbitrarily complex behaviours to be ordered and executed within a DES, on the other hand they make it difficult to model events that last in time, e.g. to simulate devices with limited computational power on which some complex task takes a not negligible amount of time to get completed. To support something similar, two reactions should be defined: one to trigger the start of the computation, and another to actually run it and complete.

5.4 Concentrations and incarnations

The low expressive power of the classical concentration is probably the hardest challenge to tackle when trying to extend a chemical-born meta model towards a richer world, where data items can be complex structures and not just simple numbers. We have found no trivial solution

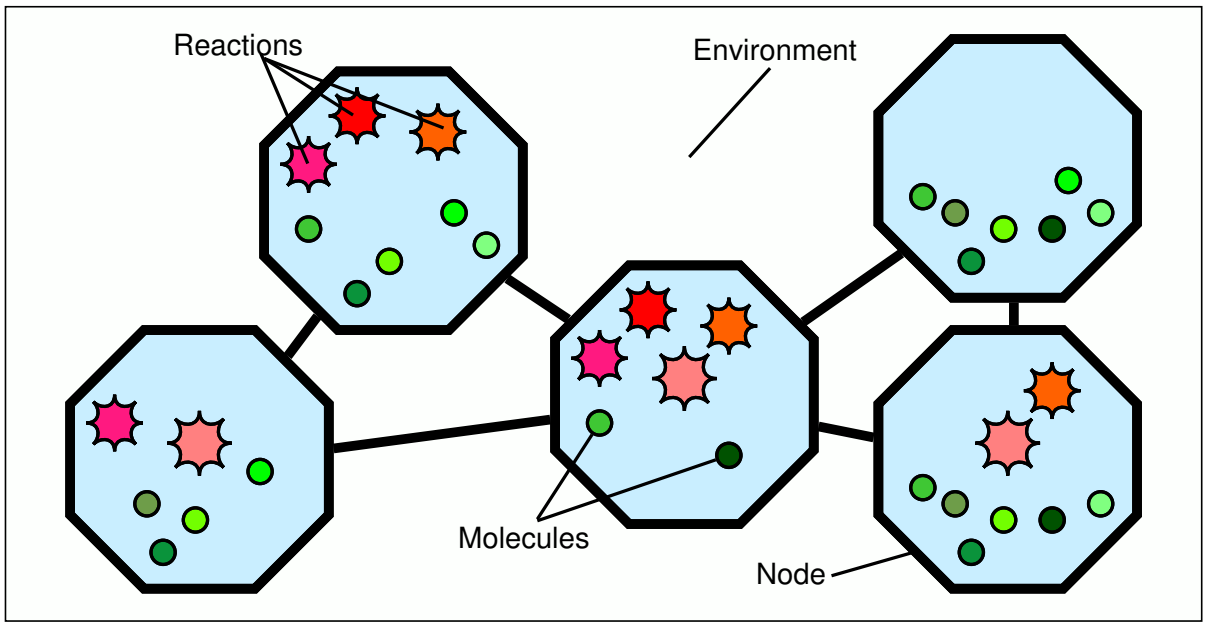

Figure 5.2: Meta-meta model: it features a possibly continuous space embedding a linking rule and containing nodes. Each node is programmed with a set of reactions and contains a set of structured molecules.

to this issue; instead, we propose to make the definition "concentration" depend on the actual meta-model: let "concentration" be "the data items agents can manipulate". Besides the trivial example of chemistry, where data items are integer numbers, let's consider a distributed tuple spaces model: in this case, the molecules would be tuples, and the concentration would be defined as "the set of tuples matching a certain tuple". Clearly, such flexibility comes with a cost: since the conditions and actions operating on different concentrations dramatically change their behaviour, for any possible class of data items the meta-model must be instanced with a proper set of conditions and actions that can act on such "concentration". We call this set of concentration-specific instances of conditions and actions an "incarnation". This sort of model inheritance justifies our double "meta" level: the model described is a meta-meta-model, an incarnation is a meta-model, and, finally, a specific scenario instancing one of such meta-models is a model.

5.5 Final remarks

A pictorial representation of the proposed meta-meta-model is shown in Figure 5.2. In this vision of the world, an environment is a multi dimensional space, continuous or discrete, which is able to contain nodes and which is responsible of linking them following a rule. The environment

may or not allow nodes to move. Nodes are entities which can be programmed with a set of reactions (internally made of conditions and actions), possibly changing over time. They also contain molecules, each one equipped with a data structure generalising on the concept of "concentration".

The specific definition of concentration, along with a proper set of conditions and actions, composes a so-called incarnation, or meta-model.

6 Architecture

This chapter will discuss architectural details of the framework: how simulations can be written, how the architecture is conceived, implementation details and performance of the engine with respect to an existing ABM simulator. We named the toolchain "ALCHEMIST" after its "extended chemistry" nature, we will use this name from now on.

The whole framework has been designed to be fully modular and extensible. The whole engine or parts of it can be re-implemented without touching anything in the model, and on the other hand the model can be extended and modified without affecting the engine. The only requirement is that the outermost model APIs remain unchanged.

It is important to note that there is no restriction about the kind of data structure representing the concentration, which can in fact be used to model structured information: by defining a new kind of structure for the concentration, it is possible to incarnate the simulator for different specific uses. For example, by assessing that the concentration is an integer number, representing the number of molecules currently present in a node, ALCHEMIST returns to its original nature: a stochastic simulator for chemistry, with addition of mobile compartments and possibly structured environments. A more complex example can be the definition of concentration as a tuple set, and the definition of molecule as tuple template. If we adopt this vision, ALCHEMIST can be a simulator for a network of tuple spaces. Each time a new definition of concentration and molecule is made, a new "incarnation" of ALCHEMIST is automatically defined. For each incarnation, a set of specific actions, conditions, reactions and nodes can be defined, and all the entities already defined for a more generic concentration type can be reused. An incarnation, namely a meta-model instance of the meta-meta model of ALCHEMIST is required for the modeller to be able to write a simulation.

6.1 Language chain

As shown in Figure 6.1, the simulations are written in a specific XML language containing a complete description of environment and reactions. This code is interpreted in order to produce an actual instance of an environment: once it is created, no further interpretation is needed in order to run the simulation. This XML code is not meant to be directly exploited by users: the XML format itself is not exactly human friendly, and XML file is often considerably big (a

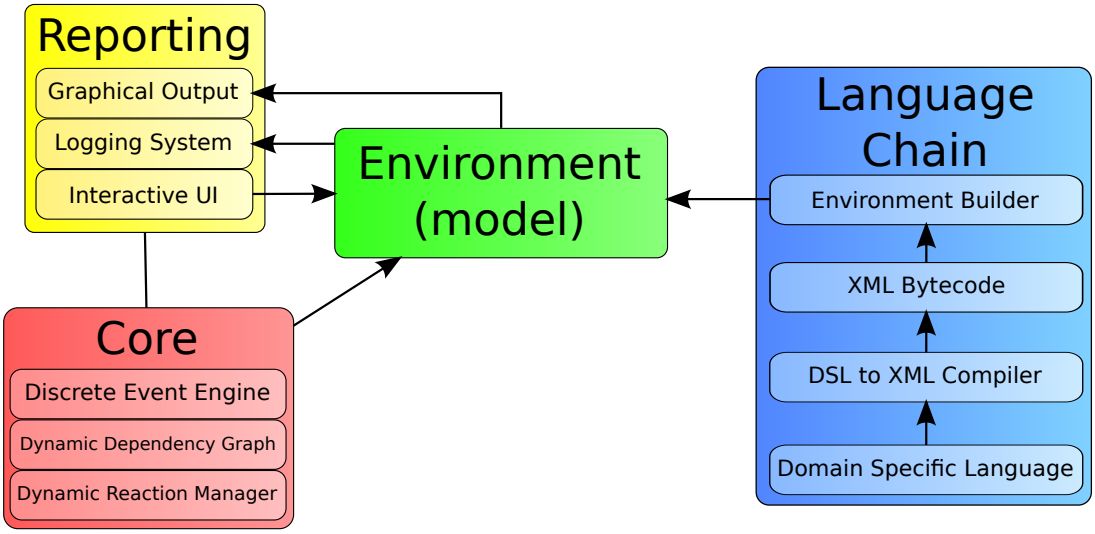

Figure 6.1: ALCHEMIST architecture. Elements drawn with continuous lines indicates components common for every scenario and already developed, those with dotted lines are extension-specific components which have to be developed with a specific incarnation in mind.

few megabytes are considered to be normal) However, it is a very standard way of describing environments in a machine-friendly format. The idea behind this choice is that ALCHEMIST is flexible enough to be used in various contexts, each one requiring a personalised language and a different instantiation of the meta-model. It is up to the extensor to write a translation module from a personalised domain specific language to the ALCHEMIST XML. Of course, it is also possible to code the simulation behaviour directly with Java, although this way exposes the user to many more low level details and it is consequently strongly discouraged.

6.2 Implementation details

The framework was developed from scratch using Java. Being performance a critical issue for simulators, we compared some common languages in order to evaluate their performance level. Surprisingly, Java performance are at same level of compiled languages such as C/C++ [BSB+03, ORAI11]. The Java language was consequently chosen because of the excellent trade off between performance, easy portability and maintainability of the code, and the high-level support for concurrent programming at the language level. Along with the very common Apache Commons Mathematics Library[1], we relied on the COLT Java library[2] provided us the mathematical functions we needed. In particular, it offers a fast and reliable random number

[1] http://commons.apache.org/proper/commons-math/
[2] `http://acs.lbl.gov/software/colt/`

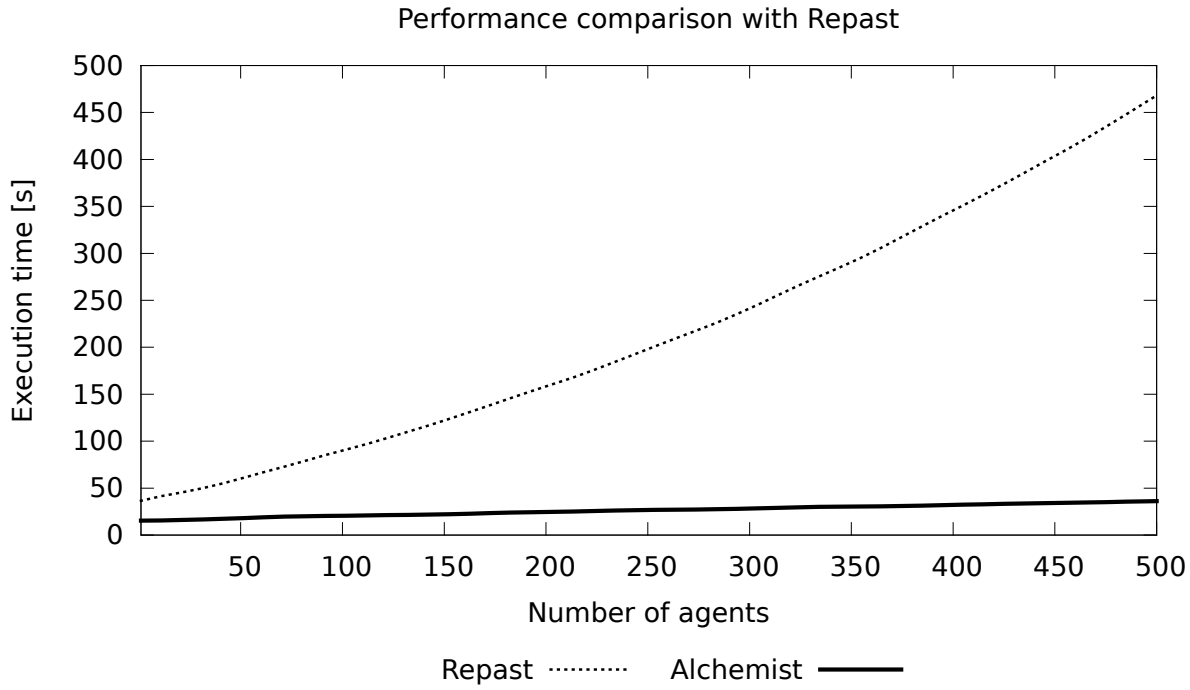

Figure 6.2: Chart showing the performance scaling of ALCHEMIST

generation algorithm, the so called Mersenne Twister [MN98]. ALCHEMIST also depends on GNU Trove[3], which provides high speed collections with very low memory impact.

ALCHEMIST is still actively developed, and we expect to soon release a stable distribution with three built-in incarnations, along with three domain specific languages. Though still under development, it is freely available[4], released under the terms of GNU General Public License 3.0[5] with linking exception.

6.3 Performance

It is possible to evaluate and compare the performances of ALCHEMIST with respect to some known MABS. We exemplify it considering Repast, which we used to developed an alternative simulation for the case study presented in Chapter 15[6]. There are some important facts that deserve discussion here.

[3] `http://trove.starlight-systems.com/`

[4] `http://alchemist.apice.unibo.it`

[5] http://www.gnu.org/copyleft/gpl.html

[6] The source code of the simulation we developed is publicly available at `http://apice.unibo.it/xwiki/bin/view/Alchemist/JOS/`.

First, since there is no built-in support for stochastic simulation in Repast, we choose to collapse the last five laws of Figure 15.1 into a single code path, and the same was made for ALCHEMIST by defining a new action. In that way, we were able to avoid for this very specific case the need of a full fledged dependency graph, because there will always be exactly one **source** and one **field** per sensor, and no reactions need to be enabled or disabled. This crippled most of the effectiveness of ALCHEMIST's dependency graph, which is indeed a source of optimisation not natively existing in Repast—developing a dependency graph for stochastic simulation in Repast is out of the scope of this thesis, though it would be an interesting subject for future work. On the other hand, it would have been unfair to compare our optimised version against just a plain Gillespie's algorithm built upon Repast.

The second important point is that we choose to encode all the data in both Repast and ALCHEMIST as an array of double values instead of real tuples. For ALCHEMIST, this meant that we developed a new incarnation for the precise scope of this performance test. The choice of encoding data that way made things faster (no matching required), but also less general. This was done because the SAPERE incarnation of ALCHEMIST, which allows us to write the laws as in Figure 15.1, requires a matching system that is not easily portable to Repast.

We choose the configuration of Figure 15.3 and we run multiple simulations varying the number of agents per group, in order to evaluate how the two systems scale with the problem size. We measured the running time required to our testbed to run the simulation from the time zero to the time 100. No graphical interface were attached to the simulators while running the batch, in order to evaluate only the raw performance of the engine. The system we used was an Intel Core i5-2500 equipped with 8GB RAM, with Sabayon Linux 7 as operating system and Sun Java HotSpot™ 64-Bit Server version 1.6.0_26-b03 as Java Virtual Machine. Results are shown in Figure 6.2. Even though ALCHEMIST apparently scales linearly, it actually does not: simply, the consistent performance gap flattens the curve.

Results show the simulator built upon the ALCHEMIST framework to be at least twice faster and to scale better than the one built on Repast. Being the dependency graph optimisation cut off as explained above, the reasons for such a difference can lay on the internal scheduler of the engine or in the optimisations in the model. For the first point, we used the default scheduler of Repast Symphony, which is a binary heap implemented through a plain Java array, while our implementation relies on the algorithm and data structures already presented in Chapter 4, so there is not a substantial efficiency difference. For the latter point, a big difference in terms of performances is due to the heavy optimizations of the neighbourhoods of the default ALCHEMIST continuous environment. Since the concept of neighbourhood was part of the computational model, it was possible to adopt caching strategies in order to ensure a fast access to the neighbourhood and a quick execution of the operations on it. This is probably the component which gives the highest impact in this case, since most interactions occur among an agent and its neighbourhood.

7 Collocation in literature

7.1 ALCHEMIST as a DES

ALCHEMIST is a discrete event simulator (DES), since it combines a continuous time base with the description of system dynamics by distinguished state changes [Zei76]. Since it allows for in-simulation modifications of the environment, it can be considered to belong to the fourth generation of DESs according to [BW08]. According to Pollacia's survey [Pol89], ALCHEMIST belongs to the category of simulators featuring "internal clock", "next-event time advance", and adhering to the "Event-scheduling World View", namely, the simulator handles events and is concerned with their effect on system states.

Apart from the meta-model adopted, the main innovative aspect of ALCHEMIST with respect to the general DES approach is its ability of optimising the "compile current event list" stage [Pol89], which ALCHEMIST quickly executes incrementally by means of the management of dependencies that the adoption of a chemical-like model enables.

As far as the simulator model is concerned, instead, the class of DES more related to our approach are those commonly used to simulate biological-like systems. A recent overview of them is available [EMRU07], which takes into account: DEVS [Zei84], Petri Nets [Mur89], State Charts [Har87] and stochastic π-calculus [Pri95]. Such an overview however emphasises that all such models require a considerable effort to map biological components into abstractions of the chosen formalism: this is because none of them was specifically developed with biology or bio-inspiration in mind. As described in Chapter 5, our model is meant to overcome such limitations, since all the enhancements to the basic chemical model we support can be seen as a generalisation of the abstractions of the works presented in literature [EMRU07], and add to them the possibly of customising with much greater flexibility a simulation to the scenarios of bio-inspired computational systems.

We should finally note that the DES approach typically contrasts the use of mathematical techniques (e.g. modelling the system by differential equations). However, the possible different choice of translating a system model to ordinary or partial different equations – which can be solved numerically in a time considerably shorter than a set of stochastic simulations – is shown to be impractical in our case. This path, explored by other works in literature [MTUG09], can provide good results for dynamics that progress at more or less the same speed, and in which

abundance of species (data-items or agents in our case) is so high that it can be relaxed to real-valued variables [EMRU07]. This is not the case for many scenarios even in system biology [Cow00, UP05], not to mention scenarios like pervasive computing where reaction rates do not change continuously, and where the effect on an even small number of agents can be key to a given system evolution.

7.2 ALCHEMIST as a MABS

Even though chemical-inspired, the meta-meta-model described in Chapter 5 holds evident relationships with multi-agent-based simulation (MABS) approaches.

According to [BMV09a], agent-based platforms for simulations can be split in three categories: general purpose frameworks with specific languages (such as NetLogo), general purpose frameworks based on general purpose programming languages (such as Repast [NHCV07]) and frameworks which provide an higher level linguistic support, often targeted to a very specific application (e.g. [WBH06]). Each approach has clearly its own advantages and weaknesses. Usually, the more general purpose is the language, the wider is the set of possible scenarios, and the wider is also the gap between the language and the simulated model. This means that an higher level language allows the user to create and tune its simulations in an easier way, on the other hand it often restricts the generality of the tool.

ALCHEMIST is meant to provide a set of meta-models, possibly each with its own domain specific language, still maintaining the possibility to extend or re-implement certain abstractions using the general purpose Java language (defining a new "incarnation", namely a new meta-model).

7.2.1 Advantages

There is a set of applications which are better tackled by ALCHEMIST. In particular, ALCHEMIST is suitable for all those simulation scenarios in which agents have relatively simple behaviour, and the notion of agent-based environment plays instead a fundamental role in organising and regulating the agents' activities [BV07] by both enabling local interactions among the proactive entities [HVUM07] and enforcing coordination rules [MOV09], allowing the modeller to shift her focus from the local behaviour of the single agent to a more objective vision of the whole MAS [SGJ07]. The idea of dealing with a strong notion of environment in multi-agent systems has been deeply developed in a series of work: other than its importance in the simulation context [HVUM07], at the infrastructure level [VHR+07] and at the methodological level [MOV09], there have been proposals of meta-models such as A&A [ORV08] and infrastructures such as TuCSoN [OZ99] and CArtAgO [RPV11]. A common viewpoint of all these works is that the behaviour of those passive and reactive components structuring the environment (e.g. *artifacts* in A&A) is well defined in terms of rules expressed as declarative conditions-imply-actions fashion. Accordingly, ALCHEMIST is particularly useful either in computing systems following

the chemical inspiration [VCMZ11, VZ10, ZCF+11], or for agent-oriented systems where the role of the environment is key, up to be a very dynamic part of the whole system—where network nodes (or, in biological terms, compartments) can move, new nodes can be spawn or be removed from the system, and links can appear or break at runtime as happen e.g. in pervasive computing scenarios [ZCF+11].

7.2.2 Limitations

Inevitably, as an attempt to build a hybrid between agent-based simulators and (bio)chemical simulators, some trade-offs had to be accepted, which ultimately makes ALCHEMIST less suited for certain classes of applications. In particular, a limitation is that ALCHEMIST's agent actions have to be mapped onto the concept of reaction. On the one hand, this makes it rather straightforward to create reactive memoryless agents [BMV09a], whose goal is just to perform rather easy computations resulting in the creation, deletion or modification of information items in the environment. In fact, since it is allowed to program different nodes with different reactions, it is easy to code reactive and context-dependent agents. On the other hand, there is neither out-of-the-box facility nor any high level abstraction useful to define intelligent agents [BMV09a]. The simulator is able to run them provided the user manually writes their whole behaviour as a single Java-written reaction—and properly specifies dependencies with other reactions. Although possible in principle, performing this task too frequently is likely breaking the ALCHEMIST abstraction, making the programmer losing the nice declarative approach that chemistry endorses, and possibly hampering the optimisation techniques that motivated ALCHEMIST.

8 Meta-model agnostic features

This chapter analyses the main features of ALCHEMIST that can be exploited regardless the specific "incarnation". Any of those feature, in fact, only depends on the meta-meta model, and consequently it is immediately available for being used as soon as a meta-model is instanced. Being so generic, they add a great value to the tool since they can be used with no need to specifically write any new code, besides possibly some minimal glue or new incarnation-related features.

8.1 Distributed statistical analysis

As we have seen in Section 3.6, simulation represents a key step in the engineering of pervasive systems. Being a simulation basically a sampling over an enormous probability space, obvious questions accompany these procedures:

- how reliable are the obtained values?
- How is the number of performed simulations chosen?
- How many simulations are required in order to state system properties with a certain degree of confidence?

Moreover, there is frequently a lack of decoupling between the model specification and the definition of the system's properties of interest: they are often embedded in the model, and their values are obtained via logging operations performed during the simulation process.

This section presents a tool meant to face these challenges, obtained chaining ALCHEMIST, with MULTIVESTA[1] [SV13], a recently proposed lightweight tool extending VESTA [SMA05] and PVESTA [AM11] which allows to enrich existing discrete event simulators with automated and distributed statistical analysis capabilities. The result is thus a statistical analysis tool tailored to chemical-inspired pervasive systems. The benefits obtained by chaining the simulator with MULTIVESTA are:

[1] `http://code.google.com/p/multivesta/`

1. a language (MULTIQUATEX) to compactly and cleanly express systems properties, decoupled from the model specification;

2. the automated estimation of the expected values of MULTIQUATEX expressions with respect to n independent simulations, with n large enough to respect a user-specified confidence interval;

3. an interactive plot as well as the generation of gnuplot input files to visualize the obtained results;

4. a client-server architecture to distribute simulations.

MULTIVESTA [SV13] is independent on the model specification language: it only assumes that discrete event simulations can be performed on the input model. The tool offers a clean interface to integrate existing discrete event simulators, enriching them with a property specification language, and with efficient distributed statistical analysis capabilities.

8.1.1 MULTIVESTA and MULTIQUATEX

MULTIVESTA performs a statistical (Monte Carlo based) evaluation of MULTIQUATEX expressions, allowing to query about expected values of observations performed on simulations of probabilistic models. A MULTIQUATEX expression may regard more measures of a model, in which case the same simulations are reused to estimate them, thus improving the performance of the analysis tasks. Moreover, the tool has a client-server architecture allowing to distribute the simulations on different machines. A detailed description of MULTIQUATEX and of the procedure to estimate its expressions is out of the focus of this work, the interested reader can deepen her knowledge reading [SV13, AMS06]. Before defining a MULTIQUATEX expression, it is necessary to specify the state characteristics to be observed. This model-specific step "connects" MULTIQUATEX with the model. The state observations are offered via the `rval`(*obs*) predicate which returns a real number for each observation, specified by the string parameter *obs*.

A MULTIQUATEX expression consists of a set of definitions of *parametric recursive temporal operators*, followed by a list of `eval` *clauses*. Each `eval` clause relies on the defined temporal operators, and specifies a property whose expected value must be evaluated.

In order to evaluate a MULTIQUATEX expression, MULTIVESTA performs several simulations, obtaining from each a list of samples (real numbers). One sample for each `eval` clause is obtained, thus all the queried measures are evaluated using the same simulations, improving performance. Based on the samples obtained from simulations, MULTIVESTA estimates MULTIQUATEX expressions with respect to two user-provided parameters: α and δ. Considering the case of simple expressions with just one `eval` clause, the estimations are computed as the mean value of the n samples obtained from n simulations, with n large enough to grant that the size of the $(1-\alpha)*100\%$ *Confidence Interval* (CI) is bounded by δ. In other words, if a simple MULTIQUATEX expression is estimated as $\bar{x}$, then, with probability $(1-\alpha)$, its actual expected

value belongs to the interval $[\overline{x} - \delta/2, \overline{x} + \delta/2]$. The case of expressions with multiple `eval` clauses is similar, with the note that the `eval` clauses may regard values of different orders of magnitude, and thus the user may provide a list of δ rather than just one. After having obtained a sample for every `eval` clause from a simulation, these values are used to update the means of the samples obtained from previous simulations (one mean per `eval` clause). If the CIs have been reached for every `eval` clause, the evaluation of the expression is terminated, otherwise further simulations are performed. Note that each `eval` clause may require a different number of simulations to reach the required CI. Once the CI of an `eval` clause has been reached, such `eval` clause is ignored in eventual further simulations performed for other `eval` clauses.

8.1.2 Integrating MULTIVESTA and ALCHEMIST

In order to chain the tools, some steps, have been tackled once and for all, while others are model-specific. Essentially, in order to allow the interaction with MULTIVESTA, ALCHEMIST has to fulfill two requirements:

1. the ability to advance the simulation in a step-by-step manner (which is provided by the `playSingleStepAndWait` method in `ISimulation` interface);
2. the ability to analyze the model status after each simulation step, providing measures in form of real numbers about properties of interest.

Since both MULTIVESTA and ALCHEMIST are Java-based, their interaction has been easily realized by subclassing the `NewState` class of MULTIVESTA. The obtained `AlchemistState` class is sketched in Listing 8.1, where unnecessary details are omitted. The new class contains some ALCHEMIST-specific code, providing MULTIVESTA with the simulation control and proper entry points for the analysis.

In the constructor (lines 6-10), the superclass initialization is done by a simple `super()` call. The remaining code initializes ALCHEMIST-specific parameters such as the maximum time or steps to simulate. Those are in general not required, since MULTIVESTA is able to detect when the analysis requirements have been met, and consequently stop the simulation flow. The method `setSimulatorForNewSimulation()` is depicted in lines 12-20. The method is invoked by MULTIVESTA before performing a new simulation to (re)initialize the status of the simulator, providing to it a new random seed. In lines 22-24, `performOneStepOfSimulation()` is provided: resorting to the ALCHEMIST method `playSingleStepAndWait()`, it allows MULTIVESTA to order the execution of a single simulation step.

In order to inspect the simulation state, the `rval()` method defined in lines 26-32 is invoked. The argument specifies the observations of interest. This method inspects the simulation state for all aspects common to any ALCHEMIST model, e.g., in the listing are sketched the current simulated time (line 27) and the number of performed simulation steps (line 28). Clearly, each ALCHEMIST model will have its own observations of interest.

Depending on the model at hand, it may be necessary to refine the model-independent observations exposed by `AlchemistState` with a set of model-specific ones. This can be

Listing 8.1: AlchemistState extending MULTIVESTA's NewState class

```
1  public class AlchemistState<N extends Number, D extends Number, T> extends NewState {
2    private final long maxS;
3    private final ITime maxT;
4    private ISimulation<N, D, T> sim;
5    ...
6    public AlchemistState(final ParametersForState params) throws ...{
7      super(params);
8      final StringTokenizer otherparams = new StringTokenizer(params.getOtherParameters());
9    /* Initialization of Alchemist-specific parameters
10      and execution environment resorting to otherParams */
11   }
12   ...
13   public void setSimulatorForNewSimulation(final int seed) {
14     /* Stop current simulation, create a new one. */
15     ...
16     sim.stop();
17     ...
18     env = getFreshEnvironment(seed);
19     sim = new Simulation<>(env, maxS, maxT);
20     ...
21   }
22   ...
23   public void performOneStepOfSimulation() {
24     sim.playSingleStepAndWait();
25   }
26   ...
27   public double rval(final String obs) {
28     if(obs.equals("time")) return getTime();
29     if(obs.equals("steps")) return getStep();
30     //other model-independent observations
31     ...
32     return getStateEvaluator().getVal(obs, this);
33   }
34 }
```

done by simply instantiating the `IStateEvaluator` interface provided by MULTIVESTA, constituted by one method only.

8.1.3 Example

This section discusses the analysis performed on a crowd steering scenario, similar to the one which will be presented in detail in Chapter 15. In such scenario, two groups of people are inside an indoor environment, and are steered towards two points of interest (POIs) such that their ideal trajectory intersect, creating a crowd. The steering system is implemented leveraging the SAPERE meta-model, and as such data items in this scenario are called `LSAs`. This example is not meant to provide significant scientific insights on crowd steering (thoughts about it will be provided in the remainder of this thesis), but rather to act as toy-example to demonstrate the flexibility and possibilities of the chaining of ALCHEMIST and MULTIVESTA.

The outcome of the analysis is summarized in the three charts of Figure 8.1 (obtained using the gnuplot input files provided by MULTIVESTA), showing, at the varying of the simulated

Listing 8.2: The evaluated parametric multi-expression (*MainMQ*)

```
1  people@POI(x) = if{s.rval("time") >= x}
2        then s.rval("bPOI")
3        else #people@POI(x)
4     fi;
5  people@RPOI(x) = // similar to people@POI
6  people@LPOI(x) = // similar to people@POI
7  avgConn(x) = if{s.rval("time") >= x}
8        then s.rval("degree")
9        else #avgConn(x)
10    fi;
11 LSAs(x) = if{s.rval("time") >= x}
12       then s.rval("LSAs")
13       else #LSAs(x)
14 fi;
15 eval parametric(E[people@POI(x)],E[people@RPOI(x)],
   E[people@LPOI(x)],E[avgConn(x)],E[LSAs(x)],x,0.0,1.0,50.0);
```

time, the expected values of: the number of people which have reached their point of interest (top), the average number of connections of the devices (middle), and the number of LSAs in the system (bottom).

The three charts have been obtained by evaluating *MainMQ* of Listing 8.2, having 5 parametric temporal operators (lines 1-9): `people@RPOI` and `people@LPOI` regard the number of people which have reached, respectively, the POI on the right and the one on the left, while `people@POI` counts instead how many people have reached their destination. The fourth temporal operator (`avgConn`) regards the connectivity degree of the devices. Finally, `LSAs` regards the number of LSA in the system, giving insights on the amount of memory required to sustain the system. The temporal operators are analysed at the varying of the simulated time from 0 to 50 seconds, with step 1 (line 10). Thus the analysis consisted in the estimation of the expected values of the 5 temporal operators instantiated with 50 parameters, for a total of 250 expected values.

A high degree of precision has been required: α has been set to 0.01, while the δ values (the size of the CIs) have been chosen considering the orders of magnitude of the measures: 0.5 for the instances of `people@POI(x)`, `people@RPOI(x)` and `people@LPOI(x)`; 0.05 for `avgConn(x)`; and 3 for `LSAs(x)`. To reach such a level of confidence, the tool ran approximately 2500 simulations, requiring less than a hour. The discussed confidence intervals are depicted in the aforementioned charts: the two lines drawn above and below the central lines represent the obtained CIs of the expected values, thus indicating the intervals in which the actual expected values lie with probability 0.99.

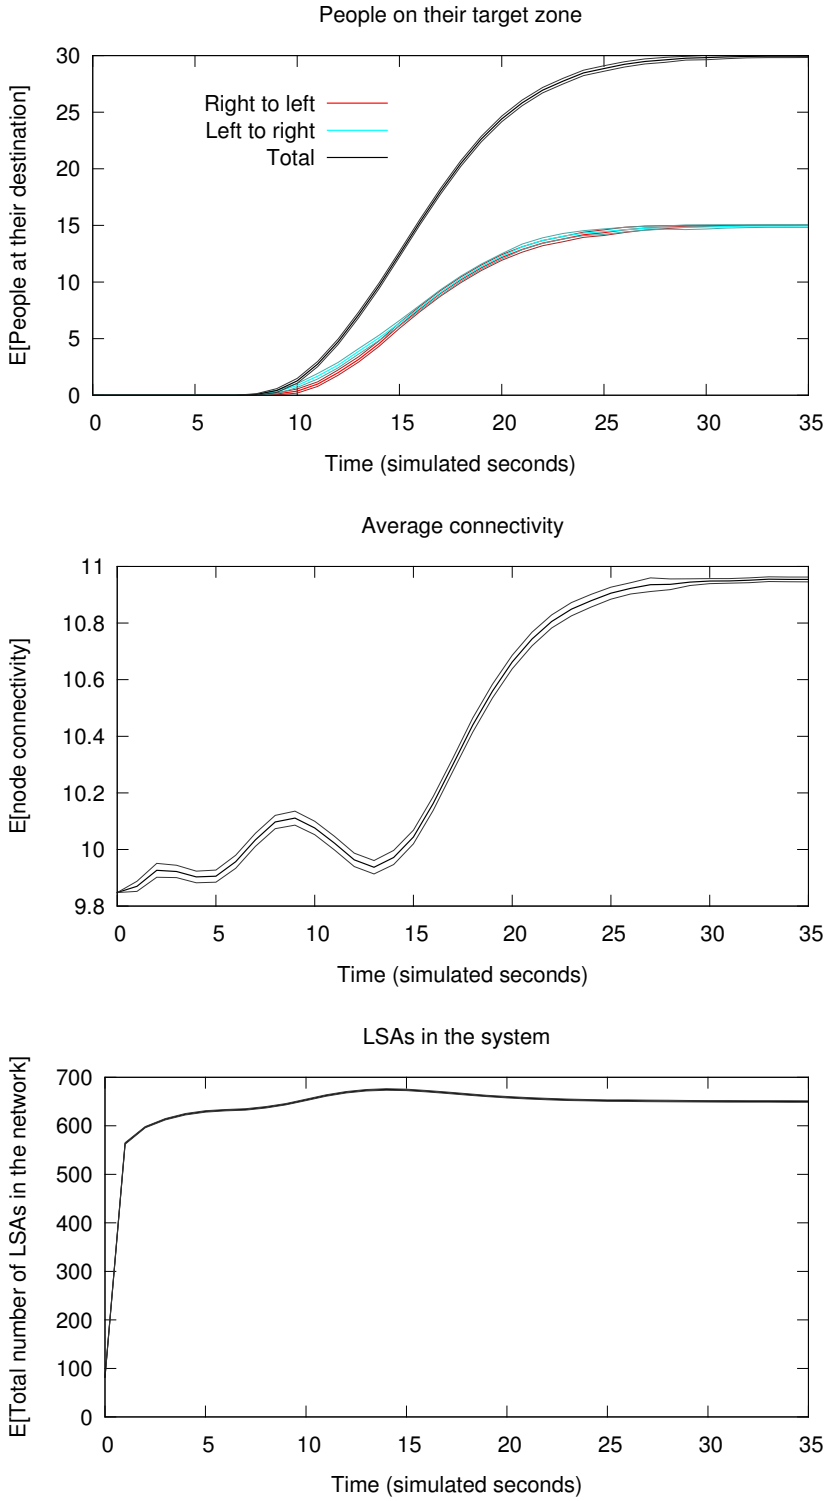

Figure 8.1: Analysis of the crowd steering scenario: (top) number of people at the POIs, (middle) average number of connections per device, (bottom) number of LSAs in the system.

The top chart regards the first 3 temporal operators: the top plot refers to the number of people in total which have reached their target POI, while the two almost over-imposed lower plots regard the number of people at the right POI and at the left POI. All the three measures under analysis produced monotonically increasing sigmoid curves. We can deduce from this behavior that there are no systemic errors in the crowd steering system, such as people with no or wrong suggested final destination, in which case we would have noticed one or more flatted zones flawing the sigmoid curve. We note also that after around 30 simulated seconds all the people have reached their target, and that it takes around 10 seconds for the fastest walkers to get to their destination. Note that, despite that the analysis has been performed for 50 simulated seconds, the charts presented in Figure 8.1 have been cut at time 35 to better show the most relevant results: in fact, the system reaches stability after this time, and no event of interest happens later. Moreover we can also notice that people going from left to right are slightly faster than the other group. This difference, due to the asymmetric positioning of people in the center of the scenario, is hardly visible and would have been impossible to spot with a lower precision analysis.

The middle chart shows the evolution in time of the average number of connections of the devices (i.e., both sensors and smartphones). Since each device is considered to be connected to all those within a range of 1.5 meters, it also gives us a hint about the crowding level. There is a noticeable peak at around 8 seconds: it is due to a high number of people approaching the central group. After that, the peak disappears when most people overtook the obstacle and are walking towards their POI. Finally, there is a growth: progressively, people reach their destination and tend to create a crowd.

The bottom chart shows the number of LSAs in the whole system, indirectly giving hints on the global memory usage. Once the system is started, there is a very quick growth, due to the gradient being spread from the sources to the whole sensors network and to the LSAs produced by the crowd-sensing. The system reaches a substantial stability after a couple of seconds. From that point on, the number of LSAs has very little variations: the system has no “memory leak”, in the sense that it does not keep on producing new LSAs without properly removing old data.

8.1.4 Performance Assessment

All the experiments of the previous section have been run on a machine equipped with four Intel® Xeon® E7540 and 64GB of RAM, running Linux 2.6.32 and Java 1.7.0_04-b20 64bit.

In order to measure the performance scaling of the tool, we ran our analysis multiple times varying the number of MULTIVESTA servers deployed. Results are summarized by the dark line of Figure 8.2, showing that with only 1 server (ALCHEMIST does not have any parallel capability by its own), the analysis required almost 17 hours, while with more than 30 servers (enabled by the chaining with MULTIVESTA) it required less than an hour. For the considered scenario, by distributing simulations we have thus obtained a more than eighteen times faster analysis.

It is worth to note that such comparison of performance is affected by the statistical nature of the analysis procedure. Intuitively, when evaluating a MULTIQUATEX expression resorting to x

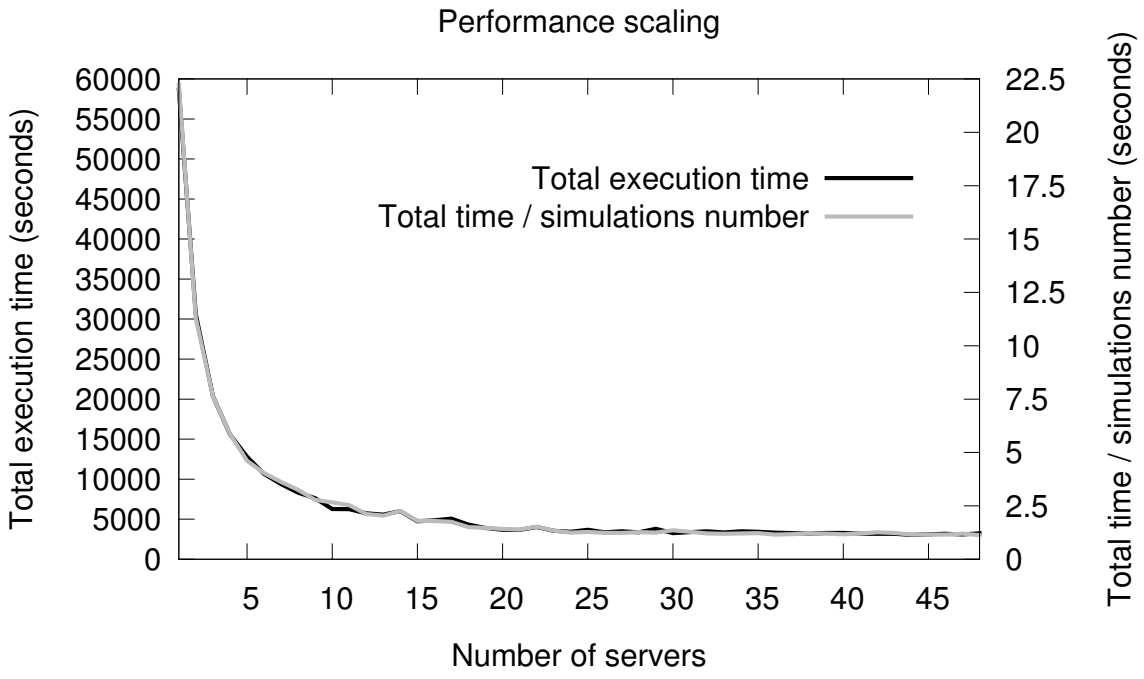

Figure 8.2: Performance scaling with the number of running servers.

or to $x+y$ servers, in both cases we initialize the first x servers with the same x seeds, generating thus the same x simulations. However, the remaining y servers are initialized with new seeds, and thus produce new simulations. Clearly, by having different simulations, we may obtain different sample variances, requiring a different number of simulations. For this reason, the bright line of Figure 8.2 depicts the time analysis normalized with respect to the number of simulations (obtained dividing the overall analysis time by the number of performed simulations). The two lines of Figure 8.2 evolve accordingly, thus confirming the great performance gain brought by the distribution of simulations.

MULTIVESTA has a further important feature which dramatically reduces the analysis time: the reuse of simulations to estimate many expected values. The reusage happens on two levels:

1. expressions can be made parametric, and thus simulations can be reused for computing the same property at the varying of a parameter (e.g., at different time steps);

2. it is possible to write multi-expressions , and thus reuse simulations to evaluate multiple properties.

Moreover, by combining these two features, it is possible to reuse simulations for multiple parametric properties (e.g. *MainMQ* of Listing 8.2). As explicated by Table 8.1 (whose reported analysis have been performed resorting to 48 servers), parametric expressions (second column)

Table 8.1: Time performance improvements reusing simulations (seconds).

	Expressions	Param. expressions	Param. multi-expressions
`people@POI`	28045.77	2567.26	n.a.
`people@RPOI`	15891.13	1114.86	n.a.
`people@LPOI`	15560.83	950.30	n.a.
`avgConn`	90111.29	1372.64	n.a.
`LSAs`	58630.45	2504.33	n.a.
Total time	208239.47	8509.39	3215.95

allowed us to save a stunning 96% of execution time with respect to the simple expressions case without reuse of simulations (first column). Moreover, the parametric multi-expression feature (third column) allowed us to further cut down this time to a third. Of course, when using parametric multi-expressions we had no way to measure the time required to measure each of the properties, since they are all evaluated in a single shot: for each simulation (sample) all the properties are evaluated, and the sequence of simulations is interrupted when all the properties are measured with the desired approximation and confidence. We can advocate that, by exploiting reuse of simulations, for the considered scenario we obtained a more than 64 times faster analysis. Clearly, the performance gains change depending on the considered scenario or properties. In fact, the more simulations are required, and the longer they take, the greater will be the advantage.

8.2 Real world maps

ALCHEMIST supports real-world environments. The base functionality lies in the possibility to load maps and simulate within them. Currently, the simulator supports OpenStreetMap[HW08], and it is able to load data in multiple OSM XML, compressed OSM XML and Protcolbuffer Binary Format (PBF). The latter is warmly recommended for both performance and map file size. Various websites provide ready-to-use extracts of the world map in PBF format for cities and whole regions, and arbitrarily sized extracts in OSM XML format can be obtained via public web API.

Once the map has been loaded, ALCHEMIST offers the possibility to enrich the simulations with the maps data, and in particular offers various ways to move nodes within the environment taking into account the characteristics of the map.

The first feature offered is about the initial node displacing. When adding a node to the environment, in fact, it is possible not to displace the node in the exact position indicated, but in the nearest street point. This comes in particularly handy in order to realise the network of static devices: the simulator can be fed with a grid of devices, and it automatically modifies the positions of each node in order to displace it on a street.

(a) GPS trace (b) Map-based navigation (c) Mixed mode

Figure 8.3: Three snapshots showing the different navigation modes available in ALCHEMIST. In Figure 8.3(a) there is a GPS trace consisting in four points. If not specified to behave otherwise, ALCHEMIST computes the average speed between two points, and moves the node along a straight line with constant speed. This, depending on the precision of the trace, may generate paths that cross over buildings. In Figure 8.3(b), the route between the start and end point is computed using the built-in navigator. This could lead the nodes towards paths different from those described by the route. In Figure 8.3(c), the routing subsystem is used to refine the quality of the GPS traces: ALCHEMIST will move the node with constant speed between the two points, using map based routing when necessary to improve precision between two consecutive trace points.

The second useful feature is the possibility to rely on the map data to compute routes, as in Figure 8.3(b). ALCHEMIST in fact ships with a module based on GraphHopper [2] which provides the simulator the ability to compute routes between two points of the map. This feature is mainly used to steer nodes correctly along the map, following the allowed ways. It is possible to use such feature also specifying different types of vehicle. Currently, pedestrians, bikes and cars are supported, and can be used together in the same simulation, and even within the same node (for e.g. simulating a pedestrian taking her car, driving and then walking again). Another usage of this system is, for instance, the possibility to use the route distance or expected route time as data items when performing internal calculations, e.g. to compute a spatial gradient where the distance is not the classic euclidean distance but rather the distance computed by the routing subsystem.

The simulator also allows for loading GPS traces. The traces must be in a personalised binary format, fortunately, however, this format is easy to generate. In fact, it is just a Java object stream containing a List of "IGPSPoints", namely a simple structure with latitude, longitude and time. A converter from JSON to such format is also available in the simulator distribution. Just as the previously mentioned routing system, the GPS traces can be used to move nodes along the scenario, reproducing existing paths, as in Figure 8.3(a).

It is also possible to use a combo of the two techniques: often, the GPS traces could be a bit rough with respect to desired grain of the simulation. In these cases, the navigation subsystem can be used to compute the route between two GPS points, with the assumption that the user followed the paths allowed in the map. In this way, it is possible to arbitrarily refine the grain of a GPS trace without the disadvantage deriving from a simpler interpolation, namely the possible transit over physical obstacles. Such mixed navigation mode is depicted in Figure 8.3(c).

Obviously, it would be rather hard to understand what is going on on the simulation without proper rendering support. In this sense, ALCHEMIST automatically detects when a real-world environment is being used, and renders a map relying on MapsForge [3].

[2] `http://graphhopper.com/`

[3] `https://code.google.com/p/mapsforge/`

9 Chemical meta-model

The chemical incarnation of ALCHEMIST was the first and simpler incarnation developed. As the reader may expect, it defines a molecule as a token, and its concentration as a numeric value. This way, ALCHEMIST comes back to its ancestral chemical nature, though retaining some of its advanced features, namely the ability of dealing with multiple, mobile compartments situated in a possibly complex environments whose number varies with time. Such abilities are not particularly interesting in a purely chemical environment, but can be key to model more complex phenomena, such as precise intracellular models or multi-cellular models, where locality and mobility become relevant.

In this chapter, we will present an application of this meta model that demonstrates the potential of this tool even outside the core applications for which it was designed.

9.1 Example scenario: Morphogenesis

Development of multicellular organisms begins with a single cell – the fertilised egg, or zygote. The egg cell is always asymmetric, *i.e.*, the distribution of maternal proteins inside the cell is not uniform. After fertilisation it divides mitotically to produce all the cells of the body. The resulting blob of cells starts the process of differentiation which is initially caused by the presence of *maternal factors* located in specific areas of the organism's egg. The ordered spatial organisation of this diversity is then caused by the interactions among cells. These can be direct or mediated by specific diffusing proteins called *morphogens*. After that, tissues have been created, and the formation of organs begins so as to originate the final shape of the organism [AJL+02, Gil06].

The main issues of Developmental Biology, nowadays only partially solved, are:

- which are the processes and the genetic mechanisms that control cellular (or nuclear) duplication;
- which are the processes and the genetic mechanisms by which cells differentiate;
- how it is possible that cellular differentiation is spatially organised.

It is clear that the macroscopic emergent results of pattern and shape formation originates from the microscopic mechanisms of intra-cellular reactions and environmental morphogen

diffusion. But how these mechanisms coordinate is still under investigation through experimental techniques and theoretical models.

A well known example of multicellular development is given by *Drosophila Melanogaster*, whose spatial organisation results in the segmented pattern of gene expression shown in Figure 9.1. To show the spatial-temporal evolution of the pattern, the organism has been object of several modelling attempts, most of which focus their investigation in gene interactions and protein diffusion mechanisms, bounding the window of analysis at a period of development that does not massively involve nuclear or cellular divisions. To cite few of them, Reinitz and Sharp [RS95] model the change of protein concentration as an Ordinary Differential Equation depending on the process of gene regulation, diffusion over a discretised space and decay. Mitoses are also modelled as a synchronous atomic process that suspends the synthesis of protein and causes a doubled number of cell nuclei. Gursky et al. [GJK$^+$04] present a continuous mathematical model based on a set of coupled non linear reaction-diffusion Partial Differential Equations that describe how protein concentrations change over time and space as a result of the mechanisms of protein synthesis and degradation, gene inhibition and activation, protein diffusion. They replace embryo's cellular structure with a continuum and do not model nuclear divisions. Montagna and Viroli [MV10] describe a stochastic and discrete version of Gursky's work based on a set of interacting compartments inside which chemical reactions implementing the gene regulatory network can occur. It does not count for changes in the network topology due to compartment movement and mitosis. These models run over a simulation window that start from cleavage cycle 11 to cleavage cycle 14. Lecca et al. [LIDP10] simulate the gradient formation of the Drosophila's *bicoid* protein, mainly caused by the morphogen diffusion, with an innovative stochastic model of reaction-diffusion systems implemented into a Gillespie-like stochastic simulation algorithm. They do not model neither gene interactions nor cellular / nuclear divisions.

To capture a bigger window of embryo development we built on top of ALCHEMIST's chemical meta-model a model of *Drosophila* development so as to capture both nuclear division and molecular processes such as morphogen diffusion and gene regulatory networks.

9.1.1 Development of Drosophila Melanogaster

Drosophila is one of the best known multicellular organism. The egg of *Drosophila* is already polarised by differently localised mRNA molecules which are called *maternal effects*. The early nuclear divisions are synchronous and fast (about every 8 minutes): the first nine divisions generate a set of nuclei forming the *syncytial blastoderm*. All the dividing nuclei share a common cytoplasm, and material can diffuse throughout the embryo. After other four nuclear divisions, during the fourteenth nuclear division, plasma membranes grow to enclose each nucleus, converting the syncytial blastoderm into a *cellular blastoderm* consisting of about 6000 separate cells. After the first ten divisions the time required to complete each of the next four divisions becomes progressively longer: cycle 13 takes for instance 25 minutes. During cycle 14 cells conclude their mitosis in rather different time: some cells take 75 minutes, other 175 minutes to complete this cycle. The rate of division is then constant in the first hours of

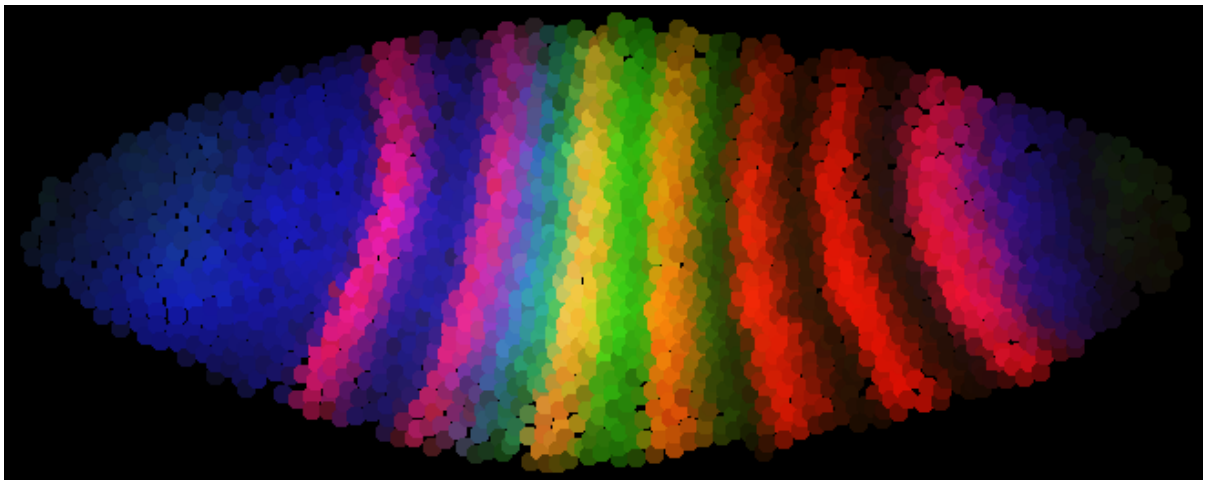

Figure 9.1: The pair-rule gene *even-skipped* (red) together with *hb* (blue) and *Kr* (green) in *Drosophila* embryo at the cleavage cycle 14A temporal class 8. Reconstructed image from [PPSR09]. Embryo name: ba3.

development (9.05 min^{-1}), then decreases until a low value (0.2 min^{-1}). The transcription of RNA massively begins during cleavage cycle 14 so that the embryo enters in the mid-blastula transition.
The studies on the genetics at the basis of the segmentation in the anterior-posterior body plan identified a hierarchy of genes that controls segment determination: maternal effects, gap genes, pair-rule genes and segment polarity genes. In Figure 9.2 an illustrative subset of this hierarchy is shown —in the images of this chapter we hereafter refer at the anterior pole as their left side, and the posterior as their right side. Maternal effect genes are the building blocks of the anterior-posterior pattern. The most important are *bicoid* (bcd) – forming an anterior-to-posterior gradient – and *caudal* (cad)—forming a posterior-to anterior gradient. The mRNAs of such genes are placed in different regions of the egg and initiate the hierarchy of transcription, driving the expression of *gap genes*, which are the first zygotic genes to be expressed. The first is *hunchback* (hb), that appears early in the embryo development such that sometimes is classified among maternal genes. The basic others are *Krüppel* (Kr), *knirps* (kni) and *giant* (gt). These genes are expressed in specific and partially overlapping domains. Differing combinations and concentrations of the gap gene proteins then regulate the expression of downstream targets, *i.e.*, the *pair-rule genes*, which divide the embryo into a striped pattern of seven segments. The most important pair-rule genes are *even-skipped* (eve) and *fushi-tarazu* (ftz). The pair-rule gene proteins activate the transcription of the *segment polarity genes*, whose protein products specify 14 parasegments that are closely related to the final anatomical segments [AJL+02].

We developed a model that reproduce the process of *Drosophila* development from the fertilised egg until the pattern formation of the *gap genes* during cleavage cycle 14. The whole embryo is modelled as a big cell where nuclei grow and move over a 2D continuous environment filled with diffusing proteins. The mechanisms we explicitly model and we reproduce with

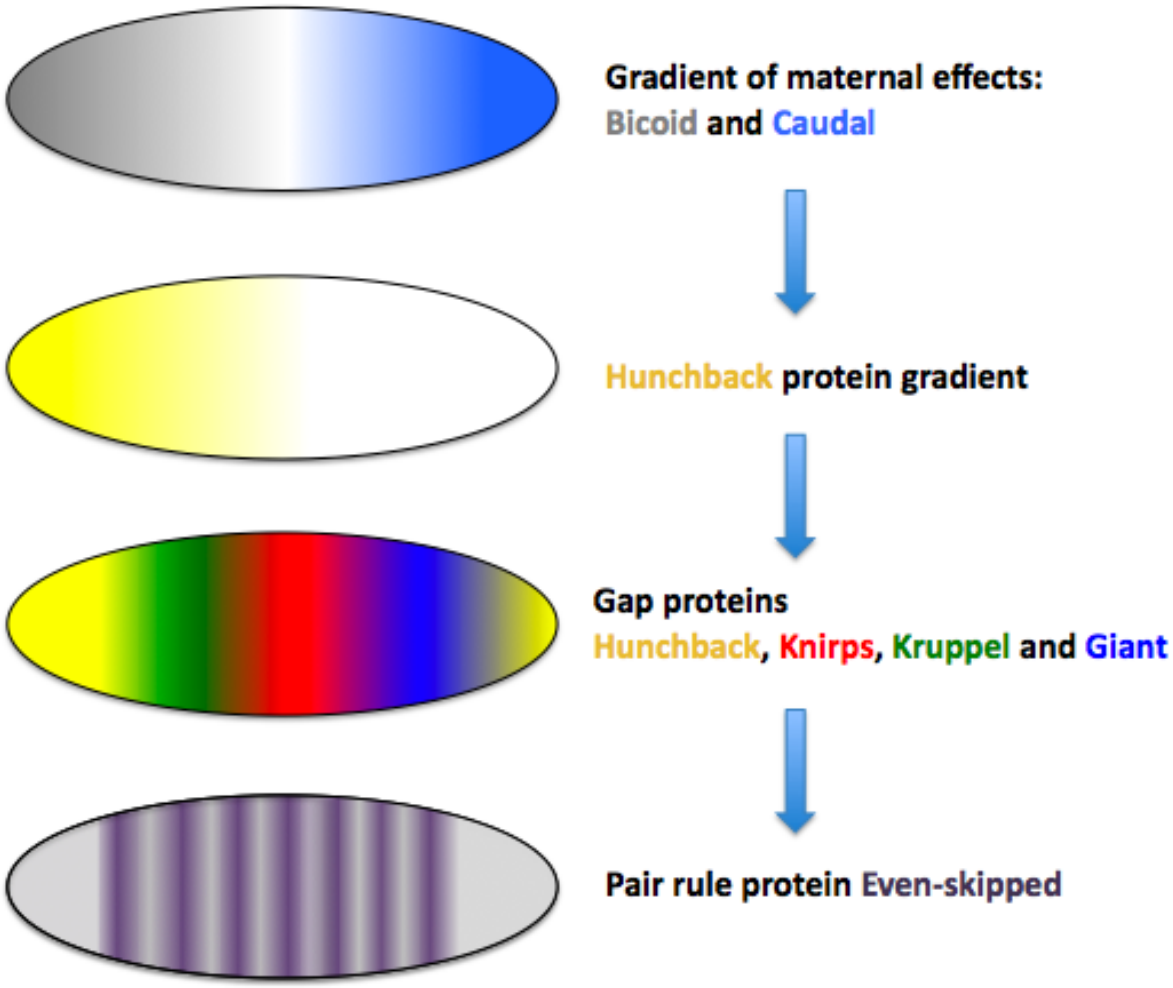

Figure 9.2: Hierarchy of genes establishing the anterior-posterior body plan

simulation are:

- nuclear divisions;
- nuclear migration;
- morphogen diffusion;
- gene interactions.

We model the process of nuclear divisions as a single chemical like reaction, whose precondition is given by the maximum number of other nuclei in the neighbourhood and whose product is a new nucleus. The new nucleus is created close to the dividing one in a casual direction and owns half of its molecular content. The rate of nuclear division is determined according to the rate observed in the real system. Since experimental data show a certain synchronisation among dividing cells, this phenomena has been modelled through a non-Markovian reaction, relying instead on a "drifting" Dirac Comb, namely on a distribution whose events happen every increasing time interval.

Movement of nuclei is based on biomechanical forces of repulsion among neighbouring nuclei. They are in fact constrained to remain within the membrane-delimited area so as to filling pretty homogeneously the available space. If two nuclei are closer than a distance given as a parameter, a new position for them is computed. The nearer two nuclei are, the stronger is the repulsion.

We support diffusion from / to nuclei and inside the environment along the x and y axis. For this purpose environment is discretised into a grid of locations. Molecules move from one nucleus into a neighbour location of the environment picked up probabilistically among the whole neighbourhood, or, the other way round, from one location of the environment into a neighbouring nucleus. Diffusion among locations of the environment follows the same law, implementing a Brownian motion.

In Figure 9.3 it is shown the network of interactions among maternal effectors and gap genes we considered. Another gap gene, *tailess* (tll), also appears as input of the network whose regulation is not clear [PJRG06] and we do not represent in the model.

9.1.2 Simulation in ALCHEMIST

Simulations results are shown in Figure 9.4. They are evaluated observing the time evolution of the compartment number and of the gene expression pattern. The time evolution of cells number is compared with data we have from literature and described above. Simulations results for gap genes expression are evaluated according to experimental data available online in the FlyEx database [SKK+08][1] and reported in Figure 9.5. They provide only qualitative information and their evaluation is based on the expression area of the different genes.

[1] `http://urchin.spbcas.ru/flyex/` – last visited on December 2014

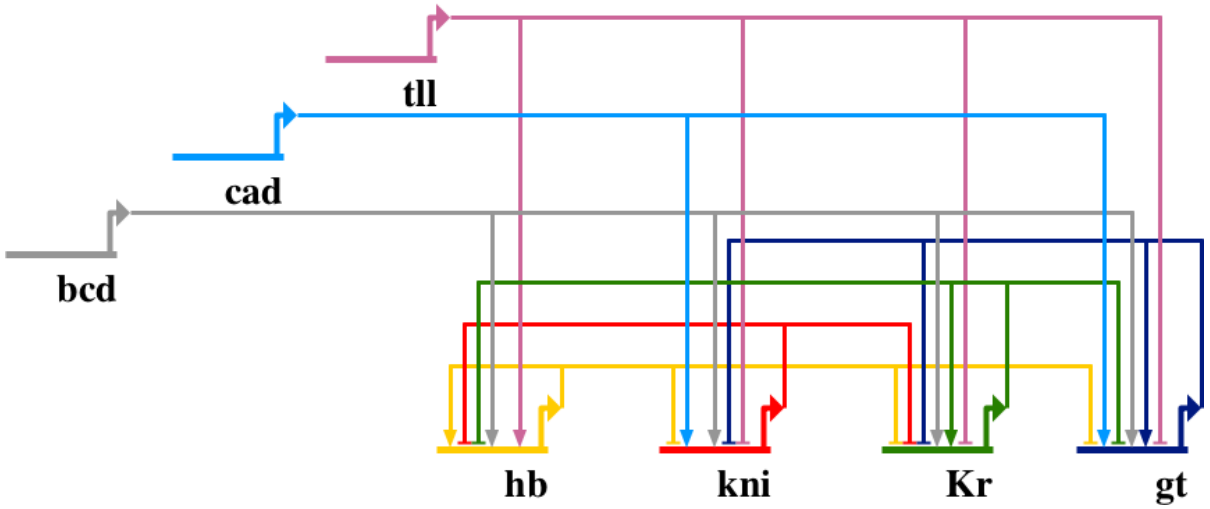

Figure 9.3: Gene regulatory relationships [PJRG06, RPJ96, GJK+04]. The diagram is realised with BioTapestry [LDB09]. The type of link is pointed out by its shape: links with arrowhead are enhancers, while the repressor links have a foot.

The snapshots in Figure 9.4 show one side of the embryo along the anterior-posterior (A-P) axis where almost half of the nuclei are located. The cell number is shown in the label on top of each snapshot. The horizontal axis represents the A-P position and is labelled as % of embryo length.

The first snapshot shows the initial condition with only one nucleus and the egg polarised by maternal effects localised in the extreme pole: *bcd* on the left and *cad* on the right. During the first minutes of simulation only nuclear divisions occur so as to fill the whole maternal cell at the end of cleavage cycle 9 with around 250 nuclei (half of the total 500), as shown in the second snapshot. Finally in the third snapshot it is reproduced the expression pattern of the four gap genes, whose spatial organisation is compared with experimental data of Figure 9.5. Cells are coloured of yellow, red, blue and green if, in order, they express *hb*, *kni*, *gt* and *Kr*, and their size is proportional to the protein concentration. Gene *hb* is massively expressed in the anterior half of the embryo and a small segment appears in the extreme left: simulations correctly reproduce the main expression domain of *hb* while, even if its expression is observable in the posterior pole, it does not form a clear segment. Gene *kni* is mainly expressed between 60% and 80% of the embryo length, either in real embryo and in simulations, as well as gene *Kr*, which is expressed between 40% and 60% of the A-P axis. The expression of *gt* is finally observable between 10% and 30% and 80% and 90%. The qualitative results presented here are not sufficient for observing the expression in the anterior half of the embryo as soon as it totally overlaps the *hb* expression, while the posterior segment clearly appears.

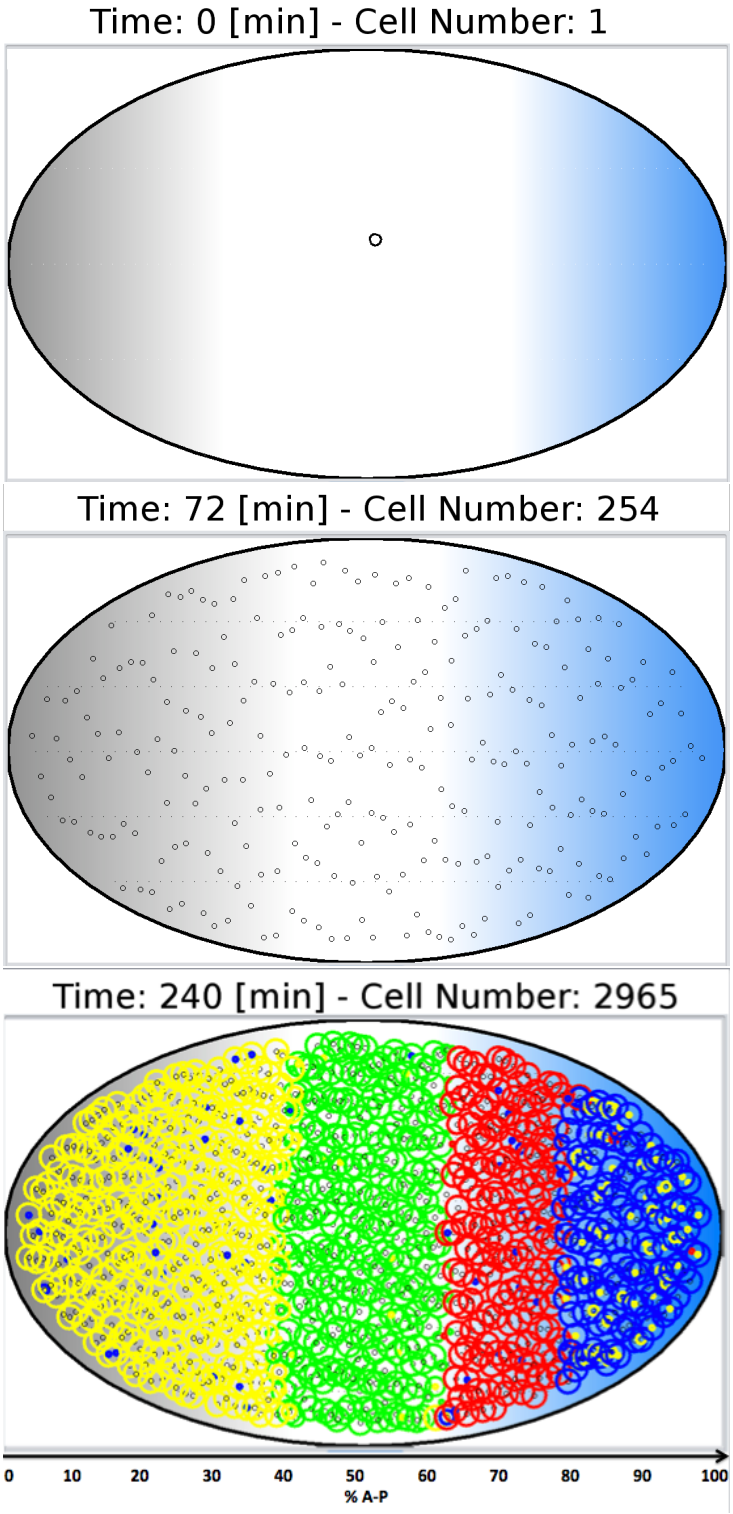

Figure 9.4: Simulation results for the four gap genes *hb* (yellow), *kni* (red), *gt* (blue), *Kr* (green) at a simulation time equivalent to the eighth time step of cleavage cycle 14A

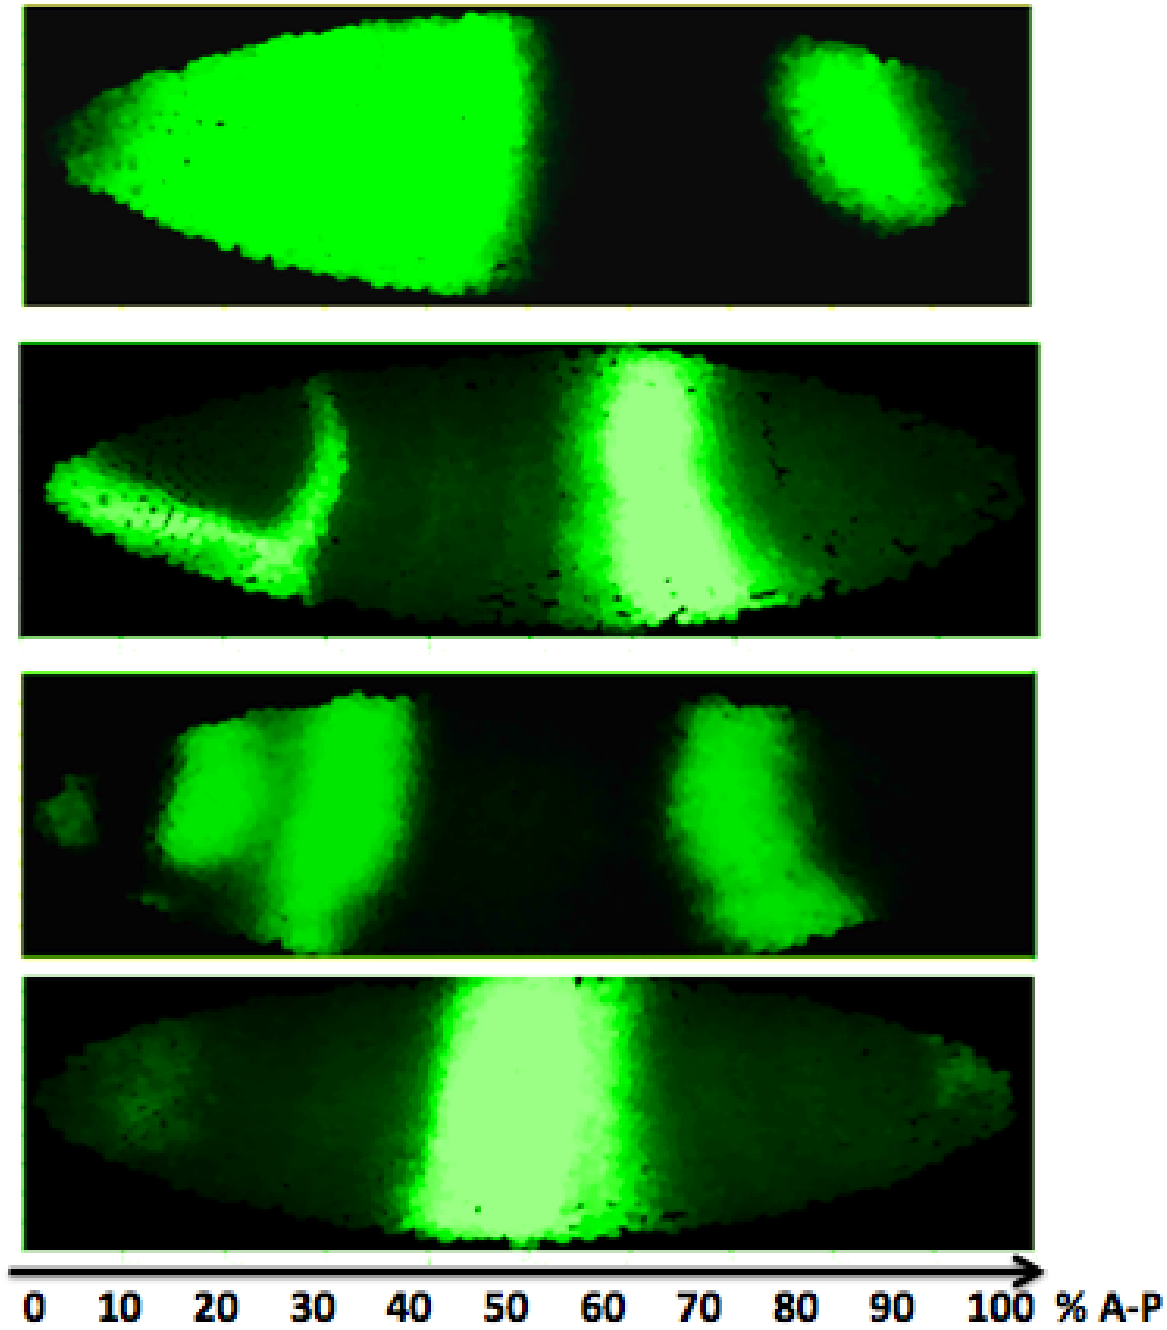

Figure 9.5: The experimental data for the expression of (from the top) *hb*, *kni*, *gt*, *Kr* at the eighth time step of cleavage cycle 14A [SKK$^+$08], ©Maria Samsonova and John Reinitz

10 SAPERE meta-model

In the previous chapter we have described the simplest incarnation of ALCHEMIST, that basically brings the model back to its chemical origins. In this chapter we will instead focus on a more complex incarnation, meant to provide support for the abstractions on which SAPERE (see Section 3.4.1) relies upon. This chapter will present no usage example, and the reason is that most of the experiments executed in Part III are run on ALCHEMIST, and rely on this incarnation.

10.1 Molecules and concentrations

In the SAPERE meta-model, a molecule is a LSA template, namely it can contain both variables and ground terms. Concentration, instead, is the collection of instanced LSAs matching some template.

To better understand this concept, a parallel with the previously presented chemical incarnation may be of use. In a multi-compartment chemical model each compartment contains some collection of chemical species, and given any of them it is possible to compute its concentration, which is the number of molecules of such chemical in the compartment. It is also possible to know the concentration of a chemical which is not present inside such compartment: the concentration would be zero. Similarly, in a multi-node SAPERE model, each node contains some collection of instanced (without unbound variables) LSAs. For each possible LSA, regardless the fact that it contains variables, we can compute the "concentration" of the template in the node: it is the collection of ground LSAs stored inside such node and matching the template. It is also possible to compute the concentration of a template that have no matches inside some node: the concentration would be an empty collection.

10.2 Reactions and eco-laws

The mapping in Section 10.1 allows the meta-meta model to be instanced in such a way that it describes a network of tuple spaces. Still, we do not have programmability, while in the SAPERE way the space is programmed with so called "eco-laws".

Programmability is obtained by properly instancing reactions, actions and conditions of the meta-meta model. For what concerns conditions, the most simple is "is an LSA matching some template present in at least one copy inside this node?". This suffices for most scenarios, but we also provide a "is an LSA matching some template present in at least one copy inside this neighborhood?". More elaborated conditions would be obviously possible, but most of them, at least for the studies that have been carried on to date, are actually completely generic and do not require to be specified inside the incarnation. For instance, "am I connected to at least some number of devices?" is totally generic.

Actions are considerably more complicated. The basic ones are rather simple: inject a new LSA locally, send an LSA to a neighbour, send an LSA to all neighbours. Actions, in the SAPERE incarnation, bust have access to the bindings performed by conditions in order to instance their templates before operating on them. Many more actions are conceivable: as example, if we wanted to implement chemotaxis, we need the ID of the node holding the highest gradient value to be used as destination for our LSA diffusion. The quickest and cleanest way to do so is to write a new action. The specific case of chemotaxis is already solved by a dedicated action, but many more may be required by other applications.

The concept or reaction, along with time distribution, already works just fine. There are only two relevant additions: the reaction must be able to read the possible variable substitutions from its conditions to pass them to its action, and it must change the way it computes its speed. In fact, since the reaction is responsible for providing a speed parameter to the time distribution, in case multiple matches are possible it must consider itself as a "super-reaction" and pass as parameter the sum of the parameters computed for each possible substitution. Once the reaction should execute, it must choose the substitution to be passed to its actions probabilistically according to the specific "speed" of each single match.

Optimised gradient

The gradient pattern is of paramount importance, being the foundation of many other patterns. For this reason, ALCHEMIST's SAPERE incarnation provides a facility, in form of dedicated reaction, that implements gradient in an efficient way, working around all the machinery in place for the other conditions and actions and with solid optimisations.

10.3 SAPERE domain specific language

The SAPERE meta model is quite mature, since it was by far the most widely used and heavily developed incarnation in ALCHEMIST. As a consequence, it also includes its own domain specific language, built relying on Xtext[1].

Such language allow the modeller to specify multiple elements, and for each of them it allows to load personalised extensions of the basic interfaces, by explicitly stating which Java class

[1] `http://www.eclipse.org/Xtext/`

should be instanced, and which parameters should be used to detect and use its constructor.

Environment – By default, a plain, unlimited bidimensional space is loaded. More complex environments, such as the one described in Section 8.2 can be easily loaded by specifying them explicitly.

Linking Rule – There is no default, the modeller must explicitly say which rule should be responsible to connect nodes. Since the most frequent rule is "connect to everybody within some communication range", a teaspoon of syntactic sugar is provided for quickly specifying just the range.

Concentration – By default, a factory of collections of LSAs is used. In no cases there was need to use a personalised value.

Position – By default, bidimensional continuous positions with euclidean distance measures are used. Also, bidimensional discrete positions with Manhattan distance are provided. Particularly useful in case of in-maps experiments are positions that use latitude and longitude as coordinates, and compute distances keeping in mind the fact the the Earth is an oblate spheroid. In case a three-dimensional environment gets developed, it will need its corresponding position format to be specified.

Random seed – Allow to select a precise simulation to be run. Very useful for debugging models. By default, instructs the simulator to generate a random simulation ad each run.

Time – Normally, time is represented by a double. In some scenarios, if time is very long but very fine granularity is needed, double values may suffer precision loss. This option is meant to circumvent the problem, provided that the modeller has her own implementation of the Time interface.

LSAs – Syntactic sugar that allows the modeller to bind frequently used LSA templates to variables, in order to shorten and ease the rest of the model code

Programs – Groups of reactions

Reactions – Reactions are specified using a compact syntax if the default type is used. In such a case, they are just a list of conditions, a time distribution and a list of actions. A reaction can't be empty: it must have at least one condition or one action.

Conditions and Actions – Can be specified inside reactions. If just an LSA is written, then the default "match an LSA"/"create an LSA" conditions/actions are used. Unary operators can be prefixed to specify that the condition/action should work on neighbourhood rather than on local node. The reactions have a chemical-like semantics: if a reactant is written on the left hand side but not on the right hand side, an action that removes it will be created.

Time distribution – By default, a "as soon as possible" scheduling policy is assumed. If only a number or an equation is specified, such number is used to initialise an exponential distribution (it is its λ value).

Group of nodes – Nodes can be specified one by one with their position or in groups. In the latter case, the modeller can choose if she wants the node of the group to occupy a rectangle or a circle. Once the shape is chosen, it is possible to specify how randomly the nodes should be displaced inside: they can be displaced in any level between perfect regularity to total randomness by acting on a parameter. For each group, it is possible to specify the content of the nodes (the LSAs that will be present when the model is loaded) and their program. Contents can be injected selectively, by specifying a spatial constrain: only nodes whose position respect such constrain will get the content injected.

Scenario number – Controls the positioning of nodes, it is a sort of seed for the scenario: since the position of nodes in groups can be randomised, there is need to be able to reproduce the exact same scenario, e.g. for debugging.

10.4 Final remarks

In this part, we have presented a novel toolchain supporting the engineering of pervasive computational ecosystems. The idea behind the tool is to provide a very generic meta-meta model, on which languages and coordination models can be mapped, in an operation of model-to-model, generating a meta-model. Such meta model can then be instanced, along with a description of an environment, into a model that can be fed to the simulator engine.

We took inspiration from the stochastic simulators for chemistry, which are known for their high performance, and we extended one of the most efficient algorithms (Gibson/Bruck) up to the point that it is able to support multiple compartments, communications, situatedness and mobility.

We defined our meta-meta model as an extended version of chemistry that includes the concept of environment and communication model, and possesses considerably more generic concept of reaction. Also, we identified what a meta-model (or incarnation) should define: its own concept of "concentration", namely data type which it manipulates, and possibly a set of targeted conditions and actions.

We introduced two important tools which are incarnation-agnostic: a probabilistic model checker obtained chaining ALCHEMIST and MULTIVESTA, and the support for real-world maps. Finally, we introduced two of the available incarnations, supporting biochemistry (with an example of morphogenesis) and supporting programmable tuple spaces with the SAPERE metaphor.

Part III

Methods and patterns for self-organisation

11 A process algebra for SAPERE

This part of the thesis presents a number of contributions in terms of novel spatial patterns and case studies. The first chapter introduces a possible formalisation of SAPERE through a process algebra. Such formalisation provides a way of resolving possible ambiguities, and also guidance for the construction of simulations or middlewares. Such abstract formalisation will be concretely instanced in Chapter 12, along with a simulated example. Subsequent examples will be provided using a simpler and possibly more intuitive syntax, but they could be mapped onto the formal model presented in the following.

We formalised the SAPERE concepts in terms of a model whose evolution is specified by means of a process algebra. Not only does this formalisation act as a non-ambiguous description, it is also an executable specification and can hence serve as ground for developing simulations of self-organisation mechanisms and for deriving formal proofs of behavioural properties, like self-stabilisation [Vir13]. Few works attempted at a process algebraic formalisation of self-organisation [VPB12, VDB13], developing on top of well-known works on coordination and spatial computing [BGZ98, BB06]. However, the one presented here is – to the best of our knowledge – the first process algebra of self-organising multiagent systems capturing the clear separation of agent behaviour and self-organised interactions, hence effectively expressing the behaviour of large-scale and situated pervasive computing scenarios.

A suggestion for the reader, to better understand how the algebra works, is to make sure she understand the example in Section 11.2.3.

11.1 Model

We now detail a possible coordination model of pervasive ecosystems. It is important to note upfront that the presented coordination model can be instantiated in different ways, depending on the application at hand and on technological aspects: This affects the concrete structure of data, which is an issue orthogonal to the above aspects. A similar abstraction is rather common in the description of space-based coordination models, which often represent tuples as simple atomic elements [BGZ97].

11.1.1 LSAs

LSAs have a unique, system-wide identifier (LSA-id), needed to support a notion of identity that is key both to uniquely identify the agent that injected an LSA and to properly support a bonding mechanism based on reference and not on value/copy. We refer to the content of an LSA as its *description*, which includes all the information the agent wants to manifest to the ecosystem. Several concrete approaches to the structure of an LSA description can be taken. For the sake of exemplification, following existing works [VZSD12, VPMS12, SYD+13], we take a semantic approach based on Resource Description Framework (RDF) representation [MM04], in which an LSA is a set of multi-valued properties, or equivalently, a set of triples (subject, predicate, object) of an LSA-id, the property name, and the assigned value, expressed as literals (strings) or URIs (terms qualified by universally-accessible namespaces, to which an ontological interpretation can be given [RMCM04]). An example LSA advertising an exhibition in a smart city, to be diffused around, would be:

```
[ app:type = app:poi, app:event = app:exhibition,
  app:artist = "caravaggio" "michelangelo", ...]
```

URI `app:artist` represents a property name, assigned to the literal values `"caravaggio"` and `"michelangelo"`. Alternatively, a value can be an URI again, like `app:exhibition` that is associated with `app:event`. Namespace `app` is used here to refer to a publicly accessible ontology supporting semantic reasoning on top of those properties and values [SYD+13].

11.1.2 Agent primitives

Standard tuple-based coordination models allow agents to insert, read, and remove any tuple (by primitives `out`, `rd` and `in`) without specific access control constraints—although some models explore their orthogonal adoption, like SecSpaces [GLZ06] and coordination contexts [ORV05, RVO04]. The model we present is profoundly different, mainly due to the chemical metaphor of bonding and the manifestation/observation approach it entails, which guarantees a better modularisation and control of accesses of agent behaviour—see a deeper discussion in Section 6. In particular, the following 4 primitives are provided:

- `new(Description):Id` — This operation relates to creating a new agent manifestation in the system. It takes a description (namely, a property-value association), and correspondingly generates a new LSA and inserts it into the local LSA-space. A new identifier, associated to that LSA, is then automatically created and returned to the agent: it will be used to perform subsequent operations on the LSA. That agent will be called the *owner* of the LSA.
- `update(Id,Update)` — This operation relates to the (continuous) manifestation of agent state. It takes the identifier of an LSA owned by the current agent, and the specifica-

tion of an update to its current description (typically, property-value re-assignments), and it updates the description of such an LSA. Note that only an LSA owner can modify it.

- `observe(Id,Query):Result` — This operation relates to the observation of an agent's context, namely, of some LSA residing in the same space but possibly belonging to some other agent. It takes the identifier of an owned LSA l and a query, and searches for an LSA bonded to l that provides a valid result to the query. If one is found (nondeterministically), the query result is returned, typically in the form of a set of variable-value associations. The semantics of bonding is orthogonal to that of agent primitives: it is dictated by the structure of a specific bonding eco-law, as described in the following.
- `remove(Id)` — Finally, this operation relates to the stopping manifestation. It takes the identifier of an owned LSA, which will then be removed from the local LSA-space.

11.1.3 Eco-laws

Without any additional mechanism, the behaviour of an agent would be isolated from those of others. Following the ecosystem metaphor, we hence introduce a limited set of eco-laws (playing the role of "the laws of nature"), with the role of evolving the population of tuples over time in a way that supports agent-to-agent interaction both within the same locality, and globally by means of self-organisation. This is achieved in terms of 5 eco-laws, two dealing with local agent-to-agent interaction by the bonding (and debonding) mechanism, and the other 3 (diffusion, aggregation and decay) supporting the mechanisms needed to create self-organising distributed structures of LSAs [FMDMSM$^+$12b]. As already mentioned, our description of eco-laws abstracts from the actual structure of LSAs and from the details of pattern matching, which we consider as orthogonal application- and infrastructure-dependent aspects.

- `bonding` — This is the eco-law regulating the creation of bonds. A bond is an asymmetric reference between two LSAs residing in the same space. As already described, the existence of a bond between LSA X and Y means that the owner of X can observe the structure of Y via the `observe` primitive. This eco-law makes sure that a bond between two LSAs exists as long as their description matches according to an infrastructure-specific matching function. In the typical case, a bond is created between an LSA providing a (possibly semantic) query and an LSA positively replying to the query. By denoting with "$\rightsquigarrow Y$" the existence of a bond to Y, this eco-law has the chemical-like structure[1]:

$$X + Y \rightarrow X^{\rightsquigarrow Y} + Y$$

- `debonding` — This eco-law is dual to the previous one, and regulates destruction of a bond between X and Y provided the conditions for the bond are no longer satisfied, namely, when either X or Y are removed or changed. It has the structure:

[1] Symbol "+" is used as separator as in chemical laws.

$$X^{\rightsquigarrow Y} + Y \rightarrow X + Y$$

- `diffusion` — The diffusion eco-law is used to create a clone of an LSA to be shipped to all neighbours of the current node. This eco-law applies to an LSA X that has a particular structure (typically, given property-value associations). The LSA X' accordingly created has a new identifier along with an updated value of some special properties that can keep track of the increased distance from the source (the notion of distance can be extended to consider advanced forms of context-dependency [SYD+13]). Note that such a new LSA has no owner: it is typically used to create a distributed structure of LSAs representing the global outcome of a single LSA initially created by an agent. By denoting with notation $\overrightarrow{X'}$, an LSA X' to be spread to neighbours, we can express the structure of this eco-law as:

$$X \rightarrow X + \overrightarrow{X'}$$

- `aggregation` — As the diffusion eco-law keeps creating copies of an LSA that are shipped to neighbours, the aggregation eco-law has the goal of addressing the implied multiplicity—ensuring that a single version of an LSA coming from the same source is present in a node. In particular, this eco-law takes two compatible LSAs X and Y (typically, two coming from the same "source" LSA after an iterative diffusion process) and aggregates them into a single one X' exploiting a suitable order-independent aggregation function. It has the structure:

$$X + Y \rightarrow X'$$

- `decay` — In order to provide a temporal cleanup mechanism, useful to refresh information and the shape of spatial structures, a decay eco-law is also considered, which takes an LSA X and erases it as soon as a given condition on its content holds. Denoting ε as the empty set of LSAs, it has the structure:

$$X \rightarrow \varepsilon$$

11.2 Formalisation

In this section we provide a formalisation of the proposed coordination model as informally presented in previous section. We adopt the style that has now became a standard after a series of previous works [BGZ97, VO06, VPB12, TNN12, BLM13, ML12, LBF13]. Namely, this is presented in the form of a process algebra, in which a system state is conceived as a "soup" (or networked set of soups) where agents and LSAs float: subjective coordination is realised by agents performing coordination primitives amounting to interactions with the local space

σ	Node identifier
n, m	LSA identifier
x	Identifier variable
d, e	LSA description
Δ	Update specification
$\overline{v}$	A meta-variable over set of elements $\{v_1, \ldots, v_n\}$

Figure 11.1: Meta-variables

(producing/consuming/modifying LSAs), while objective coordination is realised by eco-laws executing as rules continuously manipulating the space on top [VC09, OZ99, ROD02]. This approach has been shown to elegantly express the semantics of the various constructs in an orthogonal way, that is, one rule of operational semantics per coordination primitive (and this will indeed be the case for our formalisation as well). Additionally, using process algebras paves the way for advanced analysis similarly to what developed in related approaches, including mathematical proofs of behavioural properties [Vir13], simulation [VPMS12, DLM05], and observational equivalence [LBF13].

11.2.1 Syntax

We let meta-variable σ range over node (or device) identifiers, n, m over LSA identifiers, x over variables to be eventually substituted by those identifiers, d, e over LSA descriptions, and Δ over update specifications. Given any of those meta-variables, we use the overline notation to indicate a meta-variable over sets of elements (e.g., $\overline{n}$ is a meta-variable over sets of LSA identifiers)—and similarly for the others. Figure 11.1 recalls these meta-variables in tabular form for the reader's convenience. Notation $|\overline{n}|$ is used to extract the number of elements in $\overline{n}$.

The syntax of the calculus is reported in Figure 11.2. A system configuration C is a multiset composition (by operator $|$) of spaces of the form $[S]_\sigma$ (S is the content and σ the identifier), and neighbourhood relationships $neigh(\sigma, \overline{\sigma})$, meaning σ has the neighbourhood $\overline{\sigma}$. A space (content) S is itself a multiset composition of agents $(\overline{n})A$ (with behaviour A and owning LSAs with identifiers $\overline{n}$), LSAs $n\langle d \rangle$ (with identifier n and a description d), LSAs to be spread to neighbours $\star\langle d \rangle$ and bonds $n \rightsquigarrow m$ (from n to m). Agent behaviour A can be of four kinds:

1. 0 is the empty agent;
2. $a;A$ is the agent executing action a and then the continuation A;
3. $A : B$ is a try-catch construct (execute A, and when/if a failure occurs, B is executed);
4. $\texttt{iterate}^B\{A\}$ is iterative (infinite) execution of A preceded by execution of current iteration B. Namely, it consumes

$$
\begin{array}{rcll}
C & ::= & 0 \mid [S]_{\sigma} \mid neigh(\sigma, \overline{\sigma}) \mid (C \mid C) & \text{Configuration} \\
S & ::= & 0 \mid (\overline{n})A \mid n\langle d \rangle \mid \star\langle d \rangle \mid n \rightsquigarrow m \mid (S \mid S) & \text{Space} \\
A,B & ::= & 0 \mid a;A \mid A : B \mid \texttt{iterate}^{B}\{A\} & \text{Agent} \\
a & ::= & x\texttt{:=new}(\Delta) \mid t\texttt{.upd}(\Delta) \mid t\texttt{.obs}(\Delta) \mid t\texttt{.rem()} & \text{Action} \\
t & ::= & n \mid x & \text{Term}
\end{array}
$$

Figure 11.2: Syntax

$$
\begin{array}{ccc}
C \mid 0 \equiv C & C \mid C' \equiv C' \mid C & (C \mid C') \mid C'' \equiv C \mid (C' \mid C'') \\
S \mid 0 \equiv S & S \mid S' \equiv S' \mid S & (S \mid S') \mid S'' \equiv S \mid (S' \mid S'') \\
0 : A \equiv 0 & \texttt{iterate}^{A}\{0\} \equiv 0 & \texttt{iterate}^{0}\{A\} \equiv \texttt{iterate}^{A}\{A\}
\end{array}
$$

Figure 11.3: Congruence table of a SAPERE algebra

Note that in the iteration construct the appendix B is not present in the surface syntax: it is initially set to 0 and will dynamically keep track of the state of current iteration.

Actions correspond to the 4 coordination primitives described in previous section: $x\texttt{:=new}(\Delta)$ creates a new LSA with description defined by specification Δ, and causing substitutiion of x with the LSA identifier; $t\texttt{.upd}(\Delta)$ updates the LSA with term identifier t by specification Δ; $t\texttt{.obs}(\Delta)$ observes an LSA bond to t using Δ as a query; and finally $t\texttt{.rem()}$ removes LSA with identifier t. In particular note that as soon as they are executed, the latter three constructs should have term identifier t be bound to an actual identifier n.

A congruence relation is also introduced in Figure 11.3, which equates terms that have to be considered syntactically equivalent. The first two lines state that parallel composition operator is actually a multiset one; the third line defines properties of try-catch (if the left-hand side is over, the whole agent process is over) and iteration (iterating the empty process gives the empty process, and when current iteration is over we spawn a new one).

11.2.2 Semantics

The semantics of the calculus is parametric in the functions shown in Figure 11.4, which are to be concretised to make the model fully executable—as we will develop in next section. The filtering function gives semantics to updates and queries over LSA descriptions as follows. First of all, given an update specification Δ and an LSA description d, then $\Delta \rhd d$ gives the description obtained by applying Δ to d—this function is assumed to be total. Second, Δ can be seen as a query: we say it matches a description d with result θ (a binding of variables) if $\Delta\theta \rhd d = d$, namely, when $\Delta\theta$ is seen as an update, it would not affect d—intuitively, e.g., an LSA of type a,

$$
\begin{array}{rl}
\Delta \triangleright d = d' & \text{Description filtering function} \\
\delta(d) = d' & \text{Diffusion function} \\
\alpha(d,e) = d' & \text{Aggregation function} \\
\varepsilon(d) & \text{Decay predicate} \\
\mu(d,e) & \text{Bond matching predicate}
\end{array}
$$

Figure 11.4: Abstracted functions

$$
\begin{array}{rrcll}
[\texttt{CGR-C}] & C_0 & \to & C_1 & \text{if } C_0' \equiv C_0,\ C_0' \to C_1',\ C_1' \equiv C_1 \\
[\texttt{NEST}] & C \mid [S]_\sigma & \to & C \mid [S']_\sigma & \text{if } S \to S' \\
[\texttt{SHIP}] & C \mid [\star\langle d\rangle]_\sigma & \to & C \mid \prod_i [n_i\langle d\rangle]_{\sigma_i} & \text{if } neigh(\sigma,\overline{\sigma}) \in S,\ |\overline{\sigma}| = |\overline{n}|,\ fresh(\overline{n})
\end{array}
$$

Figure 11.5: Rules modelling architectural aspects

$$
\begin{array}{rrcll}
[\texttt{CGR}] & S_0 & \to & S_1 & \text{if } S_0' \equiv S_0,\ S_0' \to S_1',\ S_1' \equiv S_1 \\
[\texttt{TRY}] & S \mid (\overline{n})(A:B) & \to & S' \mid (\overline{n}')(A':B) & \text{if } S \mid (\overline{n})A \to S' \mid (\overline{n}')A' \\
[\texttt{CATCH}] & S \mid (\overline{n})(A:B) & \to & S \mid (\overline{n})B & \text{if } S \mid (\overline{n})A \not\to,\ A \not\equiv 0 \\
[\texttt{ITER}] & S \mid (\overline{n})\texttt{iterate}^B\{A\} & \to & S' \mid (\overline{n}')\texttt{iterate}^{B'}\{A\} & \text{if } S \mid (\overline{n})A \to S' \mid (\overline{n}')B' \\
[\texttt{NEW}] & S \mid (\overline{n})x\texttt{:=new}(\Delta);A & \to & S \mid m\langle\Delta \triangleright 0\rangle \mid (m,\overline{n})A^{\{m/x\}} & \text{if } fresh(m) \\
[\texttt{UPD}] & S \mid n\langle d\rangle \mid (n,\overline{n})n\texttt{.upd}(\Delta);A & \to & S \mid n\langle\Delta \triangleright d\rangle \mid (n,\overline{n})A & \\
[\texttt{OBS}] & S \mid m\langle d\rangle \mid (n,\overline{n})n\texttt{.obs}(\Delta);A & \to & S \mid m\langle d\rangle \mid (n,\overline{n})A^{matcher(\Delta,d)} & \text{if } n \rightsquigarrow m \in S \\
[\texttt{REM}] & S \mid n\langle d\rangle \mid (n,\overline{n})n\texttt{.rem}();A & \to & S \mid (\overline{n})A &
\end{array}
$$

Figure 11.6: Rules modelling coordination primitives

$$
\begin{array}{rrcll}
[\texttt{DIFF}] & S \mid n\langle d\rangle & \to & S \mid n\langle d\rangle \mid \star\langle\delta(d)\rangle & \\
[\texttt{AGR}] & S \mid n\langle d\rangle \mid m\langle e\rangle & \to & S \mid n\langle\alpha(d,e)\rangle & \text{if } (m,\overline{n})A \notin S \\
[\texttt{DEC}] & S \mid n\langle d\rangle & \to & S & \text{if } \varepsilon(d),\ (n,\overline{n})A \notin S \\
[\texttt{BND}] & S \mid n\langle d\rangle \mid m\langle e\rangle & \to & S \mid n\langle d\rangle \mid m\langle e\rangle \mid n \rightsquigarrow m & \text{if } \mu(d,e),\ n \rightsquigarrow m \notin S \\
[\texttt{DBND}] & S \mid n \rightsquigarrow m & \to & S & \text{if } S \not\to S \mid n \rightsquigarrow m
\end{array}
$$

Figure 11.7: Rules modelling eco-laws

matches with a query asking whether its type is X (and gives reply a), if updating the type of that LSA to a leaves it unchanged. As a mere example in this section assume Δ and d are both lists of single-value assignments $p = v$, and that filtering $\Delta \rhd d$ applies all assignments of Δ to d, such that $(p = 1, q = 2) \rhd (q = 3, r = 4) = (p = 1, q = 2, r = 4)$. Observe that update $(q = 3, r = X)$, where X is a variable, can be seen as a query which applied to $(q = 3, r = 4)$ yields $\theta = \{X/4\}$, in that $(q = 3, r = X)\{X/4\} \rhd (q = 3, r = 4) = (q = 3, r = 4)$ as described above.

The other four functions in Figure 11.4 work with LSA descriptions, and define the data-related aspects of eco-laws as described in the following.

The operational semantics is defined as a transition system with judgement $C \to C'$, meaning that system configuration C can move to C' by a transition. The operational semantics is split in three parts: general congruence rules (Figure 11.5), rules defining agent behaviour (Figure 11.6), and finally rules defining eco-laws (Figure 11.7), described in turn.

In Figure 11.5, rule [CGR-C] is the standard rule making congruent configurations be considered equivalent, rule [NEST] allows one to recursively enter a space, and finally rule [SHIP] defines the broadcast semantics: when an LSA-to-be-spread $\star\langle d \rangle$ occurs in a space, a copy of it is sent to all neighbours (by action $\texttt{ship}(\overline{n})$, where $\overline{n}$ is a set of externally provided freshly-generated identifiers)—note that, as in other works [IPW01], we abuse the notation writing $\overline{n}$ for $n_1, \ldots, n_k$ and similarly for $\overline{\sigma}$.

Figure 11.6 provides the rules defining the internal behaviour of spaces, affecting agent behaviour. Rule [CGR] states congruence inside spaces similarly to what [CGR-C] does for system configurations. Rule [TRY] handles successful operations inside a try-catch construct: simply, in that case the left-hand side of operator ":" is allowed to proceed. Rule [CATCH] conversely handles a failure: when the left-hand side of operator ":" gets stuck (and is not empty), we simply allow the part on right (the handler) to carry on. Rule [ITER] handles the semantics of iteration: it simply allows the appendix to carry on—when it becomes 0 a new iteration will be spawn by the congruence relation as discussed above.

The subsequent 4 rules handle the 4 coordination primitives of the model. Rule [NEW] create a new LSA with fresh identifier n: such an LSA gets created with description $\Delta \rhd 0$ (namely, as specified in the argument of operation `new`), n is added to the set of LSAs owned by the agent, and the continuation is allowed to carry on after applying substitution of m to x. Rule [UPD] handles update operations: the LSA with identifier n should be in the space – let d be its description – which is moved to $\Delta \rhd d$ and the agent continuation is allowed to carry on. Rule [OBS] handles an observation operation of query Δ: it looks for a bond to m and accesses its description d, it then computes via function $matcher(\Delta, d)$ a minimal substitution θ such that $d = \Delta\theta \rhd d$, and if one exists it is applied to the continuation—if no such LSA is found instead, the agent gets stuck, and this failure could be caught as seen above. Rule [REM] models removal of an LSA, which also causes removal from the set of identifiers of LSAs owned by the agent.

Finally, Figure 11.7 defines the semantics of the 5 eco-laws, which are parametric in functions δ, α, ε, and μ. The first is the diffusion eco-law: which takes an LSA whose description is in the domain of δ, and creates an LSA-to-be spread whose description is the output of δ. The aggregation eco-law takes two LSAs and applies function α to their description: the result is

stored in one of the two, while the other is removed (provided it is not an owned LSA). The decay eco-law simply drops an LSA whose description makes predicate ε hold (again, provided it is not an owned LSA). Note that the side-conditions of the aggregation and decay eco-laws prevent an LSA's removal if an agent owns it. The bond eco-law creates a bond between two LSAs if they stay in the μ predicate relation (and if one such bond does not yet exist). Conversely, the debond eco-law drops any bond that the bond eco-law would not create.

11.2.3 Example

Whereas exemplifying eco-laws will better be done in next section while discussing the case study, it is here useful to provide a simple agent program to show how the operational semantics works. Assume the following agent process A_0 is, executed into an initially empty LSA-space:

$$x{:=}\texttt{new}(q=2);\texttt{iterate}^0\{x\,.\,\texttt{obs}(e=\texttt{true});x\,.\,\texttt{upd}(q=0):x\,.\,\texttt{upd}(q=1)\}$$

By rule [NEW] a new LSA is created, assume it has id n, and substitution $\{n/x\}$ is hence propagated in the continuation leading to:

$$(n)\texttt{iterate}^0\{n\,.\,\texttt{obs}(e=\texttt{true});n\,.\,\texttt{upd}(q=0):n\,.\,\texttt{upd}(q=1)\}\mid n\langle q=2\rangle$$

Let A be the specification inside iteration, because of congruence we have that current state is equivalento to:

$$(n)\texttt{iterate}^{n\,.\,\texttt{obs}(e=\texttt{true});n\,.\,\texttt{upd}(q=0):n\,.\,\texttt{upd}(q=1)}\{A\}\mid n\langle q=2\rangle$$

Now, the appendix is allowed to carry on by rule [ITER]; however since there's currently no bond from LSA n, observation is stuck, hence because of rule [CATCH] we actually move in one step to

$$(n)\texttt{iterate}^{n\,.\,\texttt{upd}(q=1)}\{A\}\mid n\langle q=2\rangle$$

and after execution of update by rule [UPD] to $(n)\texttt{iterate}^0\{A\}\mid n\langle q=1\rangle$ which is equivalent to:

$$(n)\texttt{iterate}^A\{A\}\mid n\langle q=1\rangle$$

Without bonds, observation keeps fail and the space of LSAs will not change. If instead at some point a new LSA $m\langle e=\texttt{true}\rangle$ is injected in this space, and a bond to it is created by the bonding eco-law, then observation succeeds, so at some point we would move to state

$$(n)\texttt{iterate}^{n\,.\,\texttt{upd}(q=0):n\,.\,\texttt{upd}(q=1)}\{A\}\mid n\langle q=1\rangle\mid m\langle e=\texttt{true}\rangle$$

and then after update to:

$$(n)\texttt{iterate}^A\{A\}\mid n\langle q=0\rangle\mid m\langle e=\texttt{true}\rangle$$

11.2.4 Well-formed configurations and properties

A configuration is considered well-formed if:

1. no two LSA-spaces have the same identifier;
2. for each space with identifier σ there is precisely one item $neigh(\sigma, \overline{\sigma})$ and all elements in $\overline{\sigma}$ have a corresponding space;
3. no two LSAs share the same LSA identifier;
4. iterations do not include removal of LSAs (`rem`);
5. for any agent $(\overline{n})A$ we have $\overline{n} \vdash A$.

The latter is a well-formedness judgment for agents, checking that variables over identifiers are properly managed and that no agent performs operations on LSAs it does not own. This judgment is inductively defined as:

$$\begin{array}{rl}
\bar{t} \vdash x\texttt{:=new}(\Delta);A & \text{if } x \notin \bar{t},\ (x,\bar{t}) \vdash A \\
(t,\bar{t}) \vdash t\texttt{.upd}(\Delta);A & \text{if } (t,\bar{t}) \vdash A \\
(t,\bar{t}) \vdash t\texttt{.obs}(\Delta);A & \text{if } (t,\bar{t}) \vdash A \\
(t,\bar{t}) \vdash t\texttt{.rem()};A & \text{if } \bar{t} \vdash A \\
\bar{t} \vdash (A:B) & \text{if } \bar{t} \vdash A,B \\
\bar{t} \vdash \texttt{iterate}^{A}\{B\} & \text{if } \bar{t} \vdash A,B
\end{array}$$

We now state two fundamental properties for the proposed calculus:

- **Subject reduction** — If a configuration C is well-formed, and $C \xrightarrow{l} C'$, then C' is well-formed. This property is key to guarantee that starting from a well-formed configuration, we never reach badly structured configurations.
- **Progress** — If a well-formed configuration $C \mid [(\overline{n})A]_\sigma$ allows no transition leading to $C' \mid [(\overline{n}')A']_\sigma$ (namely, the agent is stuck), then necessarily $A \equiv (n\texttt{.obs}(\Delta);A')$, and in σ there is no bond from n to m such that m's description d satisfies $\Delta\theta \rhd d = d$. This property states that in well-formed configurations operations never fail, except for the case of an `obs` that fails and has not been caught by the try-catch construct.

12 Pattern: Channel

We now present an example application of the process algebra proposed in Chapter 11, in which the abstract syntax in Chapter 11 is concretely instanced into RDF, and the operational semantics is applied to let the system progress.

Such instantiation is specifically targeted at emphasising its ability to tackle a number of interesting issues:

1. enact robust self-organising design patterns [FMDMSM$^+$12b];
2. code relevant classes of agents (contextualizers, combinators of spatial structures, initiators of aggregation/diffusion processes);
3. support openness by an RDF-oriented instantiation of LSAs.

12.1 A semantic-oriented instantiation

We give an RDF-oriented concrete definition of data representation, which can be useful in a large number of pervasive computing applications, including the example to come—a rigorous mapping to RDF and SPARQL can be given [VZSD12]. Namely, we provide:

- the shape of LSA descriptions (metavariable d);
- the space of update specifications (metavariable Δ);
- the description filtering function ($\Delta \rhd d$);
- the definition of eco-law-related functions $\delta, \alpha, \varepsilon, \mu$.

For the sake of conciseness, we describe them informally.

An LSA description d is a comma-separated sequence of multi-value assignments of the kind `p`=$\overline{\texttt{v}}$, where *property* `p` is a URI, $\overline{\texttt{v}}$ is a sequence of values, and a single value `v` is either a URI, a string or a variable (written starting with a question mark as in `?Var` as in RDF/SPARQL). An update specification Δ can be a variable, or the literal `empty` (meaning the LSA is entirely

empty), or a comma-separated sequence of assignments `p=`$\overline{\texttt{w}}$, where each element $\texttt{w}_i$, called a compound value, can be a value or an expression to be evaluated (written `eval(exp)` where `exp` is a string).

Filtering function defines the update behaviour as follows. First, to compute $\Delta \triangleright d$ all expressions in Δ are evaluated, leading to filtering function d'. Second, if Δ' is literally `empty` then the result of filtering function is simply an empty description. If it is instead a list of assignments, the filtering function applies all of them to d, and returns the resulting description d'—an assignment `p=`$\overline{\texttt{w}}$ completely rewrites any possible pre-existing one. As an example, if Δ is `[app:p = "1", app:q = "2"]` and d is `[app:p = "3" "4", app:r = "5"]`, then $\Delta \triangleright d$ gives `[app:p = "1", app:q = "2", app:r = "5"]`.

We finally need to give the definition of eco-law functions, as of Figure 11.4. First, predicate ε holds only for LSA descriptions having property `eco:decay` set to `"true"`, which are hence the descriptions of those LSAs that will be decayed.

Function δ takes the description of an LSA l with the assignments `eco:diff_function` $= f$ and `eco:diff_prop` $= p_1 \ldots p_n$, where each p_i is assumed to be a property of l assigned to a single-value v_i. The output is a description obtained by the input one with two changes: *(i)* properties $p_1, \ldots p_n$ are no longer assigned to their old values $v_1, \ldots, v_n$, but to the result of $f(v_1, \ldots, v_n)$ which is assumed to be a tuple of n values, and *(ii)* property `eco:decay` is set to `"true"` so that as soon as the LSA is diffused and processed remotely, it will then be decayed. Note that function f does not need to be total, hence the diffusion process can be defined so as to eventually terminate. For instance, an LSA with description

```
[app:p = "1", app:q = "10", eco:diff_prop = app:p app:q,
 eco:diff_function = fun:inc fun:dec]
```

will be diffused and remotely become

```
[app:p = "2", app:q = "9", eco:diff_prop = app:p app:q,
 eco:diff_function = fun:inc fun:dec, eco:decay = "true"]
```

Aggregation is achieved by binary function α defined as follows. It applies to a pair of descriptions for two LSAs l_1 and l_2 having both identical copies of the three assignments: `eco:aggr_function` $= f$, `eco:aggr_prop` $= \overline{p}$ and `eco:aggr_pre` $= \overline{q}$: again, properties $\overline{p}$ and $\overline{q}$ are assumed to be single-valued. Then, function f is applied to the values assigned to $\overline{p}$ in l_1 and l_2, say they are $\overline{v}$ and $\overline{v}'$) respectively, and the result $f(\overline{v}, \overline{v}')$ (a tuple of values) will be used to replace in l_1 the assignments to $\overline{p}$ while l_2 is removed. As a precondition for the aggregation to happen, l_1 and l_2 should have the same assignment for properties $\overline{q}$. For instance, the two LSAs

```
[app:p = "1", app:q = "10", eco:aggr_pre = app:q,
 eco:aggr_prop = app:p, eco:aggr_function = fun:sum]
[app:p = "2", app:q = "10", eco:aggr_pre = app:q,
 eco:aggr_prop = app:p, eco:aggr_function = fun:sum]
```

will be aggregated and become

```
[app:p = "3", app:q = "10", eco:aggr_pre = app:q,
 eco:aggr_prop = app:p, eco:aggr_function = fun:sum]
```

Finally, function μ realising bonding applies to an LSA l_1 with assignments `eco:bond_prop` $= p_1 \dots p_n$ and `eco:bond_prop` $= v_1 \dots v_n$, and to another LSA l_2 with properties $p_1, \dots, p_n$ respectively assigned to $v_1, \dots v_n$. Accordingly, eco-law [BOND] will create a bond from l_1 to l_2, that will persist until this matching holds. For instance, the LSA

```
[eco:bond_prop = app:p app:q, eco:bond_values = "2" ?Any]
```

will bond to an LSA with description

```
[app:p = "2", app:q = "10"]
```

The above examples illustrate a convention we shall use in the following: URIs with namespace `eco` are those describing properties and values whose semantics is associated to the eco-laws engine, those with `fun` refer to a library of underlying functions, and those with `app` are relative to the application domain ontology.

12.2 Realising the gradient

Assume an adaptive display infrastructure, where a huge number of displays are deployed more or less uniformly in a wide and possibly articulated area, like a smart city [MVFM$^+$13]. In a given location, called the *source*, a point-of-interest (POI) is deployed: several pedestrians around may be interested in reaching that POI according to a convenient path. A steering service with the goal of guiding those people to the POI can be setup by making each display provide a dynamic sign of the direction to take, automatically computed by the opportunistic local interactions between the deployed devices. As described in previous works [MVFM$^+$13, BBVT08, VCMZ11], this service can be provided by the so-called *gradient* self-organisation pattern [FMDMSM$^+$12b]: from the POI and by suitable diffusion and aggregation processes, a distributed data-structure is established which reifies in each node the distance (e.g., hop-by-hop) from the nearest source node. By making each display show a sign pointing towards the

```
% Source: injecting the source LSA at a POI
?SRC := new(app:src = true, app:type = app:grad, app:desc = ..)

% GradientAgent: from the source, prepares the diff/aggr process
?BND_LOC := new(eco:bond_prop = eco:location, eco:bond_value = "true")
?BND_SRC := new(eco:bond_prop = app:src, eco:bond_value = "true")
?GRAD := new();
iterate {
  ?BND_LOC.obs(eco:location = ?LOC);    % observes current location
  ?BND_SRC.obs(?SRC);           % observes whole source description
  ?GRAD.upd(?SRC);       % clones SRC description into the grad LSA
  ?GRAD.upd(app:src = "false",       % adds all required properties
    app:distance = "0", app:temp = "false", app:location = ?LOC,
    eco:aggr_pre = app:type app:temp,
    eco:aggr_prop = app:distance,
    eco:aggr_function = fun:minOnFirst,
    eco:diff_prop = app:distance app:temp app:location,
    eco:diff_function = fun:inc fun:falseToTrue ?LOC)
  :
  ?GRAD.upd(empty)    % if above observations fail, empties the LSA
}

% Contextualizer: if not blocked by a proper LSA, reifies the gradient
?BND_TEMP := new(eco:bond_prop = app:temp, eco:bond_value = true);
?BND_OBS := new(eco:bond_prop = app:block, eco:bond_value = ?B);
?GRAD := new();              % creates the LSA for the gradient LSA
iterate {
  ?BND_TEMP.obs(?TEMP);             % observe temporaneous grad LSA
  (                    % either consolidate or empties the grad LSA
   ?BND_OBS.obs(app:block = true); ?GRAD.upd(empty)
   :
   ?GRAD.upd(?TEMP); ?GRAD.upd(app:temp = false)
  )
  :
  ?GRAD.upd(empty)
}
```

Figure 12.1: Multi-agent system for creating a distributed gradient data structure

neighbour node with smaller distance value, one obtains the desired steering service. A solution to this problem can be obtained with the three agents described in Figure 12.1.

First of all, a *source* agent in each node located at the POI locally injects an LSA declaring the node as being a source for the POI gradient, with all additional domain descriptions of the POI.

Then, a *gradient* agent, deployed in each node of the network (or at least on source nodes), has the goal of observing the presence of a source LSA, and accordingly start the diffusion/aggregation process ending up with a gradient data structure. By insertion of LSAs `?BND_SRC` and `?BND_LOC`, the agent seeks to bond with a source LSA and a location LSA—namely, an LSA which we assum some location agent creates and maintains to reify current node's location. It also creates an empty LSA that will hold the gradient value at this position- It then iteratively reads current location (`?LOC`) and whole content of source LSA (`?SRC`), and correspondingly updates the gradient LSA: it copies there all assignments in the source LSA, then sets distance to 0, "temporaneous" flag to `false`, location to `?LOC`, and suitable aggregation and diffusion functions. Such functions make the LSA diffuse by increasing distance by 1 at each step, updating flag from `false` to `true`, keeping the current location upon diffusion, and making multiple copies of this LSA coming from different nodes re-aggregate, always taking the one with minimum distance[1].

Finally, a *contextualiser* agent is in charge of enacting the situation recognition process by which it may be decided that the gradient has not to be locally propagated: if a blocking LSA is present (one with assignment of `app:block` to `true`), the gradient will not be propagated further here—modelling e.g. a dynamically moving obstacle to be circumvented. After suitable requests for bonding as in the previous case, and creation of an empty LSA called `?GRAD`, it iteratively checks whether there is a temporaneous gradient LSA and if there is no "blocking" LSA: if both conditions hold, it simply copies the content of gradient LSA into `?GRAD` and moves the flag to `false` (so that diffusion/aggregation will carry on), otherwise it empties `?GRAD` entirely.

The gradient structure created by this process enjoys self-stabilisation [VD14], namely, it repairs in response to *any* change in the environment, including topological changes and re-positioning of POI.

12.3 From gradient to distributed channel

On top of a gradient structure one can set up a steering service that activates all the signs located along an optimal path from an area *S*, where interested people are initially located, to the POI *D*. Such a path should also have a given "width" to take into account the variability of people movement. The resulting spatial structure can be referred to as a (distributed) *channel*,

[1]Function `fun:minOnFirst` yields the first argument if smaller than the second, while function `fun:falseToTrue` applies only if the flag is `false` and turns it to `true`

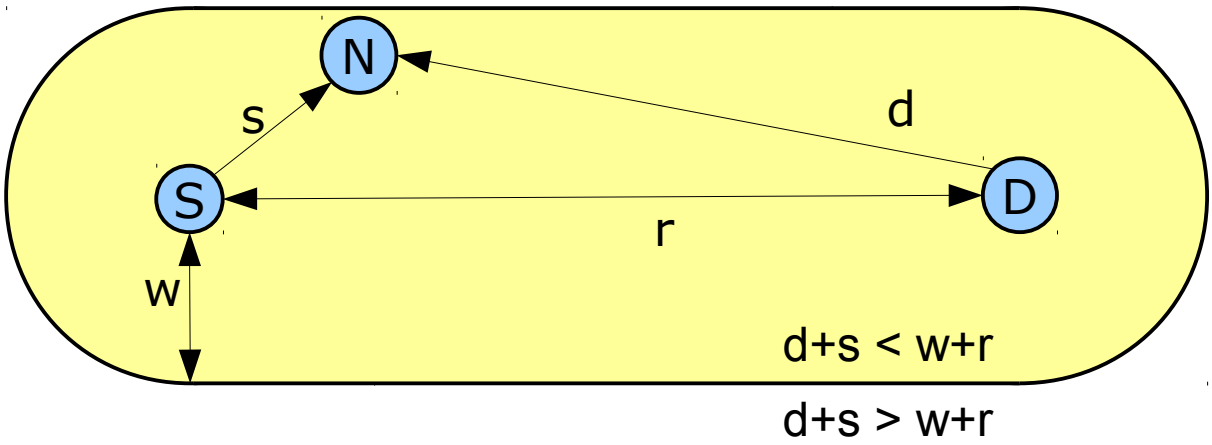

Figure 12.2: Construction of a distributed channel, based on distance between the two ends *S* and *D* (*r*), width (*w*), distance from *S* (*s*) and distance from *D* (*d*).

and should be created by self-organisation in a way that it can, again, automatically adapt to the presence of obstacles and to changes in the environment—see a preview in Figure 12.2.

A solution to this problem can be conceived following the general approach of "spatial computing" [BDU+13, Vir13, BB06], namely, functionally composing simpler gradient-based data structures. It can be coded in our process algebraic language as shown in Figure 12.3, with the behaviour of the following 3 kinds of agent.

First a *channel source* agent located at position *S* is in charge of bonding with the gradient spread from position *D*, and accordingly creates a source LSA `?SRC` which will be continuously updated to reflect the current distance `?R` from *S* and *D*, and to carry also the width value `?W`. The gradient agent defined in Figure 12.1 will be able to create a new gradient out of this source LSA, which will propagate carrying both `?R` and `?W`, in addition to the distance from *D* associated with the `app:distance` property.

Then, a *channel service* agent is able to finally create the channel LSA: it bonds and observes both gradients, and the sets a flag into a newly created (and continuously updated) LSA depending on whether the node is inside or outside the channel area, as described in Figure 12.2.

Finally, a *steering service* agent has the goal of reifying a sign towards the direction of area *D* only in nodes inside the channel area: this is achieved by retrieving the current location (`eco:location`) and the location of the neighbouring node with smallest distance value from *D* (held in `app:location`) and performing a subtraction operation—the semantics of this subtraction depends on the actual shape of coordinates used for identifying locations. The LSA it produces and maintains is representative of the direction of the sign to be shown on pervasive displays, and is intended to be observed by a display agent.

```
% ChannelSource: initiates the channel propagation
?BND_GRAD := new(eco:bond_prop = app:type, eco:bond_value = app:grad);
?SRC := new();  % the source LSA initing channel propagation
iterate {
  ?BND_GRAD.obs(app:distance = ?R);         % perceiving POI's gradient
  ?SRC.upd(app:src = true, app:type = app:rng,
           app:range = ?R, app:width = ?W)            % updating source
  :
  ?SRC.upd(clear);
}

%ChannelService: computes the final channel data structure in each node
?BND_DST := new(eco:bond_prop = app:type , eco:bond_val = app:grad);
?BND_RNG := new(eco:bond_prop = app:type , eco:bond_val = app:rng);
?CHN := new(app:type = app:chn, app:active = "false");
iterate {  % observes ?D,?S,?R,?W and computes channel activation flag
    ?BND_DST.obs(app:distance = ?D);
    ?BND_RNG.obs(app:distance = ?S, app:range = ?R, app:width = ?W);
    ?CHN.upd(app:active = eval("?S + ?D < ?R + ?W"));
    :
    ?CHN.upd(app:active = "false");
}

% SteeringService
?BND_LOC := new(eco:bond_prop = eco:location, eco:bond_value = "true")
?BND_CHN := new(eco:bond_prop = app:type app:active)
?BND_FIE := new(eco:bond_prop = app:temp, eco:bond_value = "false");
?DIR := new();
iterate {
  ?BND_CHN.obs(app:active = "true"); ?BND_LOC.obs(eco:location = ?MYLOC);
  ?BND_FIE.obs(app:location = ?DIRLOC);
  ?DIR.upd(app:direction = eval("?MYLOC - ?DIRLOC"))
  :
  ?DIR.upd(empty);
}
```

Figure 12.3: Multi-agent system for the distributed channel

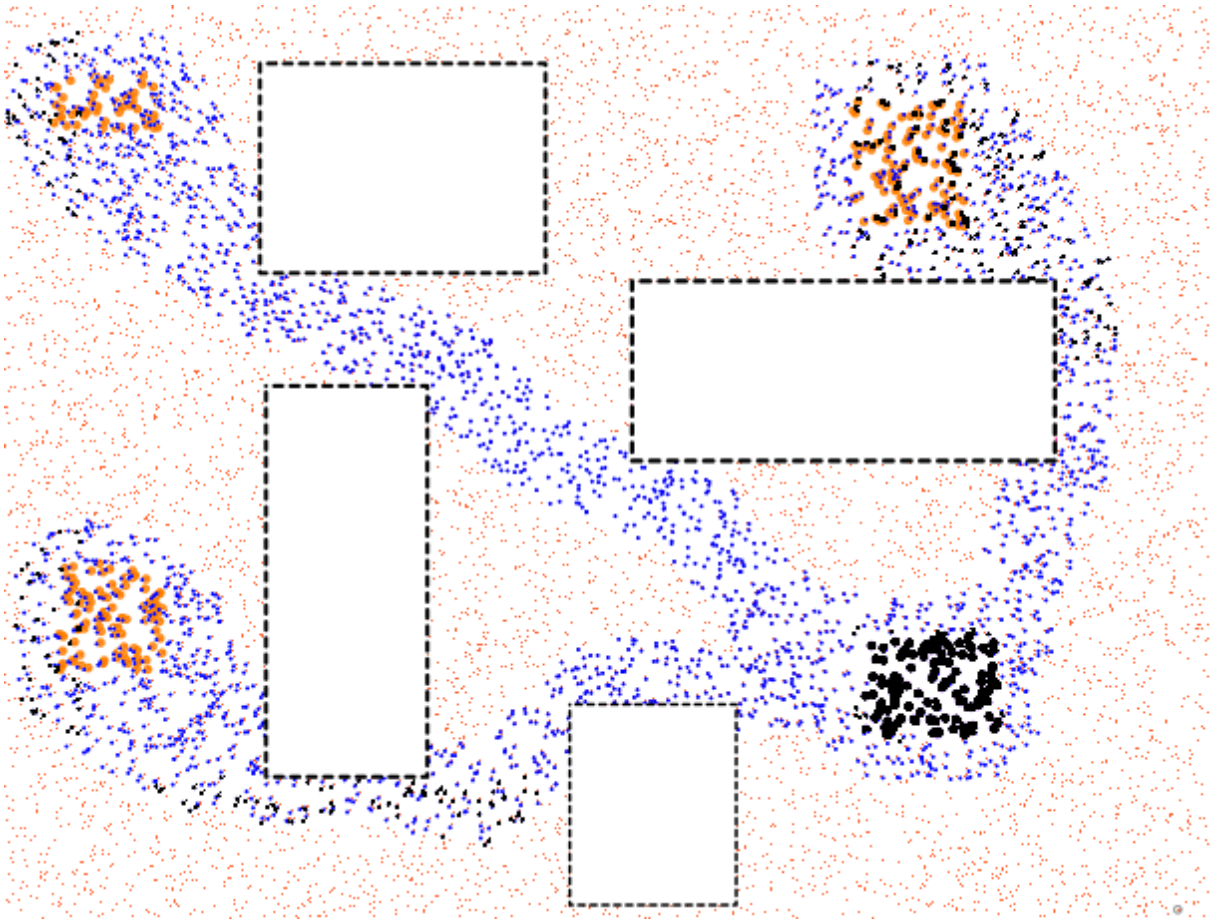

Figure 12.4: Channels (blue) route between source (orange) and destination (black), deployed in a high-density environment of 10'000 nodes (right) with blank areas acting as obstacles.

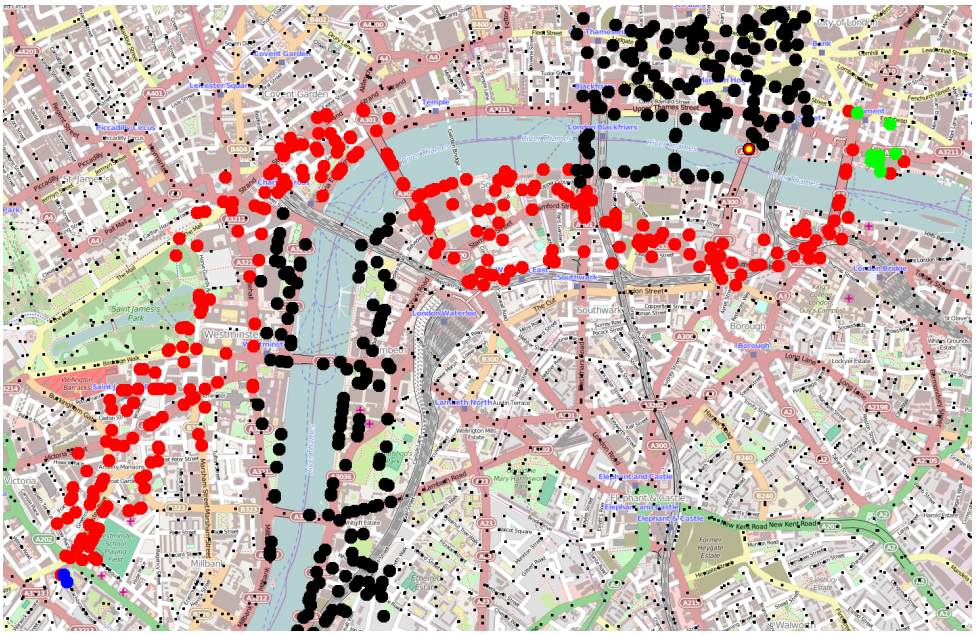

Figure 12.5: The distributed channels deployed in a smart city (London) for pedestrian steering. Red nodes represent devices inside the channel area, while black nodes correspond to “obstacle” areas.

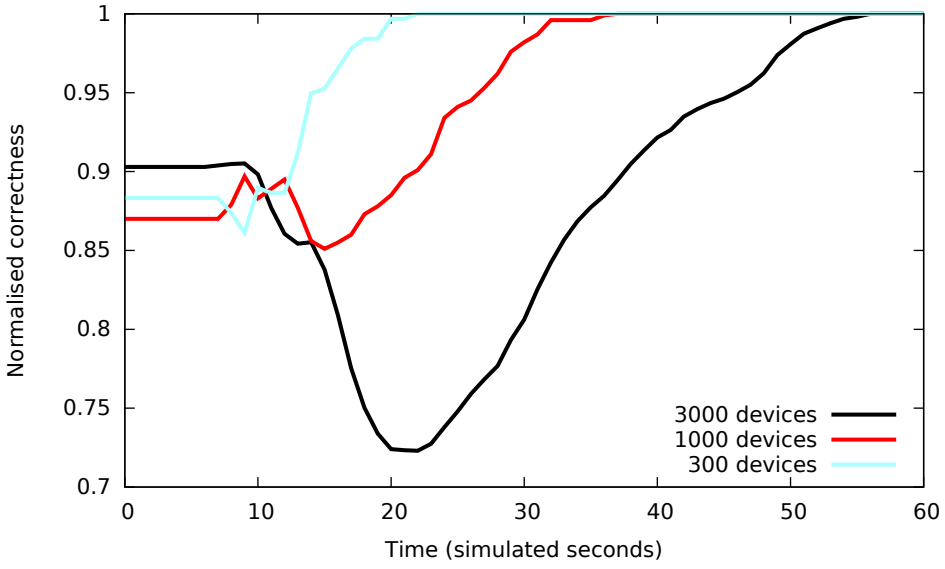

Figure 12.6: Percentage of devices correctly showing channel information over time, using topologies with different densities.

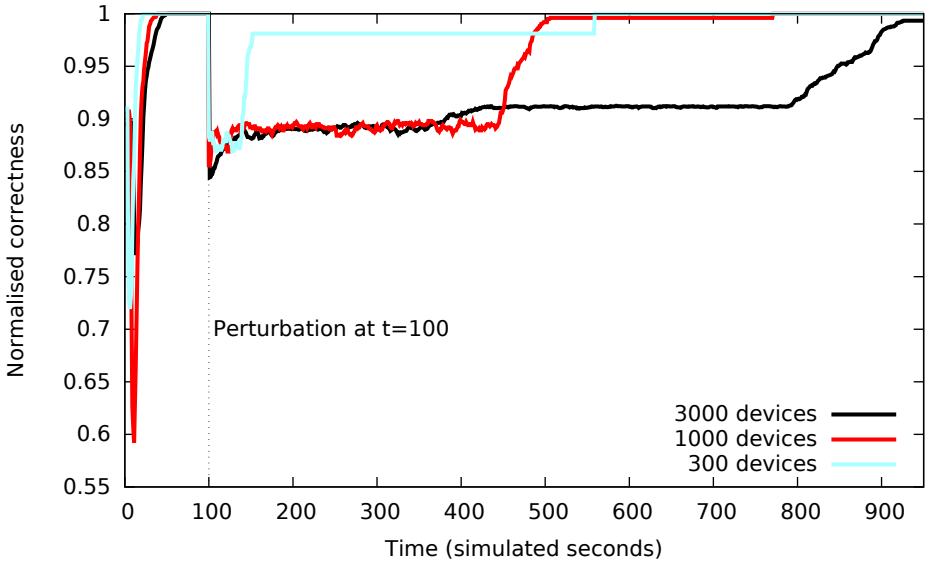

Figure 12.7: Percentage of devices correctly showing channel information in the presence of addition of an obstacle.

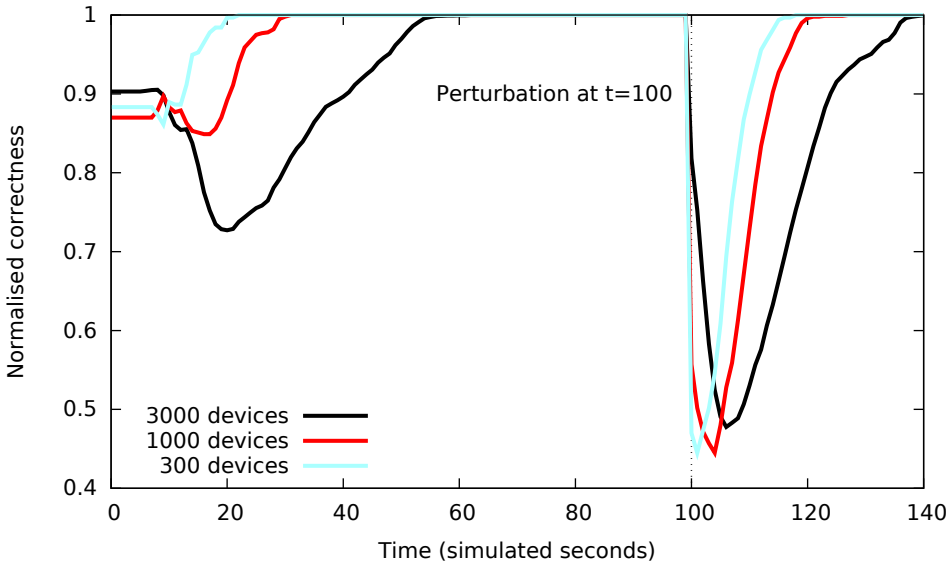

Figure 12.8: Percentage of devices correctly showing channel information in the presence of removal of an obstacle.

12.4 Simulation in ALCHEMIST

We provide a brief qualitative and quantitative account of the behaviour of the channel distributed data structure, with the goal of validating the design of agents as defined above.

We first discuss some qualitative aspects, with help of the snapshot in Figure 12.4, showing the result of a simulation of a channel structure in a high-density network, created by running an equivalent model in the Proto simulator [BB06]. First, one can note that the resulting channel data structure is highly independent of local aspects of network topology: provided that nodes are spread more or less randomly, only density affects the final shape, and in particular, the "width" of the channel[2]. Second, the way the channel is created allows one to use distributed concepts of "source" and "destination" of a channel, and even to consider the case of a channel starting from multiple sources—e.g., to guide different groups of interested people. This is because self-organisation is achieved here by the functional combination of well-engineered spatial structures, as advocated in other works [BDU+13]. Third, this structure automatically adapts to the shape of the network, properly bypassing obstacle areas. This happens since it can be shown that the diffusion-aggregation process underlying gradients enjoys self-stabilisation [VD14] (diffusion always increments distance, and aggregation is performed based on minimum values).

The presented technique can hence be usefully employed in complex environments, where it is difficult to predict the contingent situations. We consider, as an example scenario, pedestrian steering in smart cities. Figure 12.5 shows a snapshot of a simulation corresponding to a deployment over a London map. Devices were displaced with different density for each experiment, totalling 300, 1000 and 3000 devices respectively. Devices are located on streets available to pedestrians, and device connections are based on proximity. We considered as the POI Monument Street, and created a channel starting from the Royal Horticultural Society, where interested people may be initially located. Dark nodes are those in which some form of unfavourable condition is assumed to be detected, requiring dynamic re-route along different paths.

From a qualitative point of view, Figure 12.5 confirms the self-adaptive character of the distributed channel data structure, as we discussed above. From a quantitative point of view, we measured the time for this data structure to establish under different conditions. Starting from the basic scenario depicted in Figure 12.5, in Figure 12.6 we investigated the time needed to reach a full stabilisation of the correct channel data structure. It charts the percentage of nodes which reached the correct final state over time. The result is that stabilisation is always eventually reached, in an average time that grows linearly with network diameter. The speed of information diffusion is approximatively computed as the distance of nodes multiplied by the frequency at which they run the computation round of eco-laws. Figure 12.7 and Figure 12.8 study self-stabilisation in response to changes in the environment, corresponding to the addition/removal of

[2]This is because we used a hop-count-gradient to propagate distances: full independence of density can be achieved by suitably reifying estimated distances between nodes as proper LSAs, and using such information when re-aggregating LSAs—which we did not develop for the sake of simplicity.

all obstacles shown in Figure 12.5, which requires a complete (spontaneous) re-calculation of the path. These charts confirm the estimation of stabilisation time available in literature [BBVT08]. Namely, the repair time is much faster with removal of obstacles, since in this case gradient values quickly fall down to their new values (i.e., distances quickly decrease). On the other hand, with obstacles added instead, values more slowly grow to the new stable state. This happens because of the issue of slow "raising values" pointed out in previous works about gradients [BBVT08].

13
Pattern: Anticipative adaptation

We note that – to the best of our knowledge – existing techniques studied for Gradients (see e.g. [Bea09, VCMZ11, VPMS12]) tackle only present-awareness, that is, they adapt to events only at the time they occur. In this chapter , we are instead concerned with how Gradients can be equipped with *anticipative adaptation*, namely, how they can take advantage of the local availability of information about future events, once this has been provided by some component in charge of situation recognition [YDM12]. Our idea is to use such a knowledge in advance to stretch and deviate certain parts of a Gradient structure so as to promote alternative paths that could be longer, but guarantee to eventually circumvent any forbidden area so as to reach the target in shortest time. Following the self-organisation approach, this will be realised in an emergent way, relying only on local interactions and without any supervision. Example scenarios where this mechanism can be useful include traffic management, to tackle pre-scheduled road constructions, availability of historical formation about traffic jams in certain areas of a city, known timing of opening bridges or railroad crossing, and so on. It can also be used in crowd steering applications based on signs appearing in public displays [VPMS12]: e.g., in a large exhibition centre the scheduling of popular and typically crowded exhibitions can be considered with the goal of anticipating and preventing the formation of crowds.

13.1 Anticipative Gradient

We conceive a new self-organisation mechanisms built upon the above Gradient pattern to tackle anticipative adaptation, namely, ensuring that the retrieval of the gradient source proactively reacts to available information about future events, namely, impossibility of transiting across a given area—e.g. modelling interruptions of roads in traffic control applications. In this section we first motivate and analyse the problem, then elaborate a solution based on a combination of mechanisms related to the gradient pattern, and finally check the general correctness of the result by some simulation.

13.1.1 Motivations and problem analysis

A basic application of the Gradient pattern is to steer a requester from its current location to the Gradient's source location—namely, what is called Chemotaxis pattern [FMDMSM$^+$12b]. Such a requester (which we here generally refer to as an agent) can either be a mobile software agent roaming a network to find proper information, a person guided in a complex environment by signs dynamically appearing in its smartphone [PMV11], a car which is given traffic control suggestions through its navigation system, or an information item directed to the target of a communication [VCMZ11]. It is quite standard to make the Gradient field adapt to contingencies as soon as they appear: for instance, interruption of a road can reflect into a change of the network topology, leading to the automatic recalculation of the Gradient structure to identify a new path towards the target of interest. Following that line, it was studied how dynamically-formed crowded areas can be perceived as areas to be circumvented by other people roaming the environment [VPMS12].

We address the design of a different mechanism which, to the best of our knowledge, has never been considered so far: how Gradient-based self-organisation mechanisms can be extended to deal with the case in which in some node(s) there is knowledge about an event that will occur there in the future and prevent/modify the ability of an agent to transit across the node. Taking advantage of such information, one could think of a strategy to spread this information around such that – in an emergent way – the Gradient can somewhat deviate to anticipate the effect of the future event, possibly promoting a path that could be longer but guarantee to circumvent the forbidden-area eventually reaching the target in shortest time. There are several example applications for one such mechanism, tackling the availability of information about future road constructions, statistic formation of traffic jams in certain areas of a city, timing of opening bridges or railroad crossing, scheduling of popular and typically crowded exhibitions in a museum, and so on.

We aim at designing a solution that is not tight to a specific application context, but has the same level of generality of the Gradient pattern. As such we assume – in continuity with the basic hypothesis used for the Gradient pattern – that the environment is densely covered with computational nodes hosting annotations and forming a topology reflecting the structure of the environment (and hence its ability of being roamed by agents). We also assume that in each node estimated "distance" to all neighbours is available: such a distance should measure the expected average time needed for an agent to move from one location to another—here we will rely on general "time units" without any correspondence with a physical time. The underlying middleware is in charge of keeping this information lively updated, such that formation of crowded areas are automatically reflected into a higher distance of the link between two nodes, thus making a path including such a link less appealing—accordingly, we will not consider crowd information further. We hence refer to *agent-speed* as the non-uniform ratio between physical distance and estimated transiting time at each location.

In some node of the network there could be some (future-aware) situation recognition component [YDM12] in charge of intercepting a future event, namely, what happens and when.

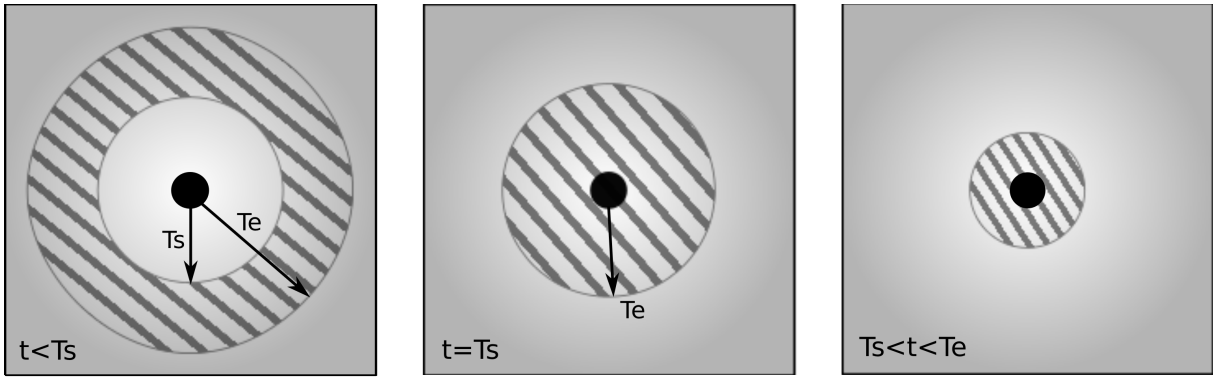

Figure 13.1: Horizon Wave pattern, with its shrinking dynamics in evidence.

We shall focus on events concerning impossibility of transiting across the node, expected in the period of time in the future in between Ts and Te ($Ts < Te$) time units from current time. Our basic strategy, detailed below, is to make information about such (possibly multiple and overlapping) events spread around, and suitably combine with a Gradient created to reach a target of interest. The result of such a combination, what we call an Anticipative Gradient, will be used by agents to navigate towards the target of interest along optimal, dynamically computed paths taking in consideration those future events. According to the self-organisation style, such mechanisms will be based only on local interactions, and will be carried on without any global supervision.

13.1.2 Spatial structures

Horizon Wave

A first component in our strategy is the advertisement of a future event (FE). We call FE area the set of nodes that will host such an event. Advertisement of the FE needs not require full-space and -time propagation, but – at a given time – it should reach only the set of nodes that might be concerned by the FE, namely, those from which an agent moving towards the FE area would actually reach it when the FE actually occurs. In a sense, the set of points in space and time that should be advertised resembles the notion of event horizon in physics, though it is here located in the past and propagates at the agent-speed. This is why we refer to this pattern as Horizon Wave (or Wave for short).

We consider for simplicity a uniform and continuous 2D environment, allowing us to better grasp the behaviour of the proposed mechanisms. There, a Wave defines a (circular) crown-like region of space that, as time passes, shrinks into the FE area at the agent-speed, until completely disappearing as shown in Figure 13.1. Agents outside this spatio-temporal data structure needs not be concerned about the FE.

This spatial structure can be obtained on top of the Gradient pattern as follows. We let

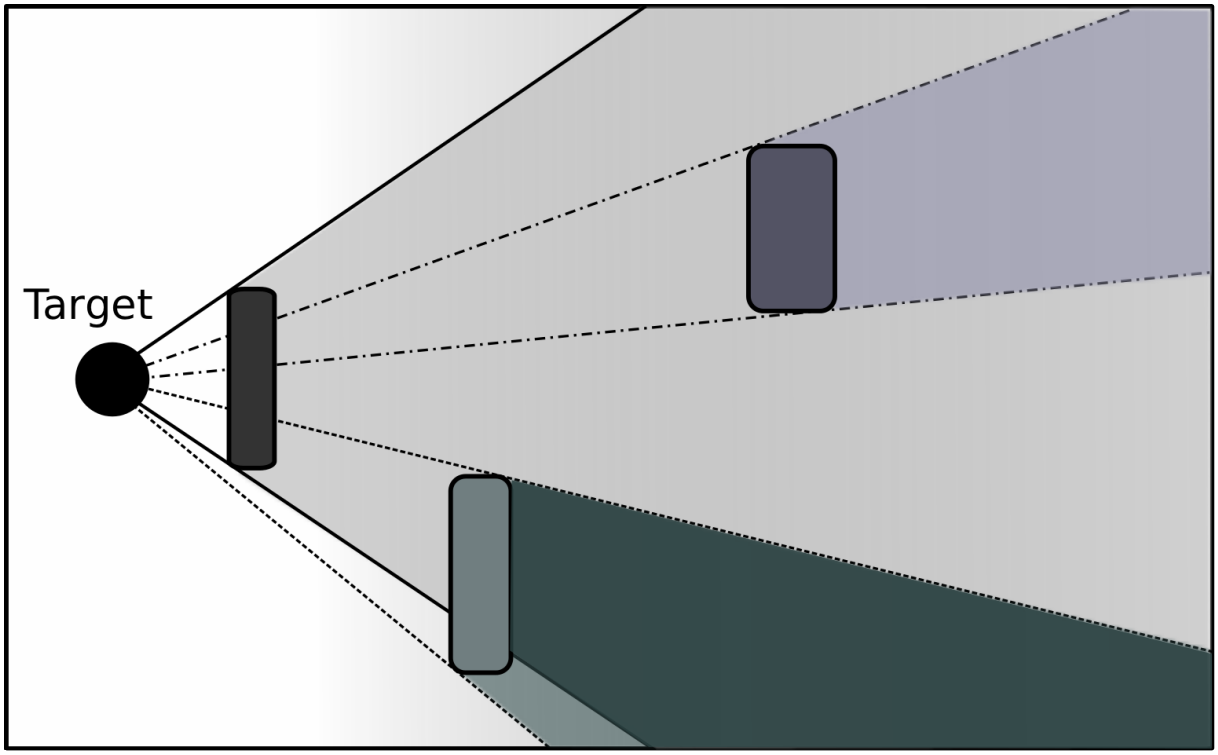

Figure 13.2: Shadow spatial structure with overlapping Future Events

the FE area be source of a Gradient, and any of its annotations (as diffused by the spreading process) carry information about Ts and Te. A node of the network is then inside the Wave if $Ts \leq D \leq Te$, where D is the node's distance from the FE area as reified by the Gradient annotation available locally.

Gradient Shadow

In order to understand whether an agent should actually care about circumventing the location of an FE, it is also important to consider its direction and target. Let's focus on the Gradient generated by a target p that an agent is directed to. We should properly tag the portion of such a Gradient from which an agent, while steered towards p, would eventually transit across an FE area. Such a portion is here called Gradient Shadow (or simply Shadow), for it recalls the shadow that the FE area would cause if it were an opaque object and the gradient source would be a source of light. Figure 13.2 emphasises that the Gradient we create out of a target of interest should be able to possibly tag different, overlapping Shadows, reflecting the existence of many FEs around.

The main strategy to create this spatial structure amounts to let Gradient propagation track, with a proper tag in the annotation, when it crosses an FE area: by successive propagation

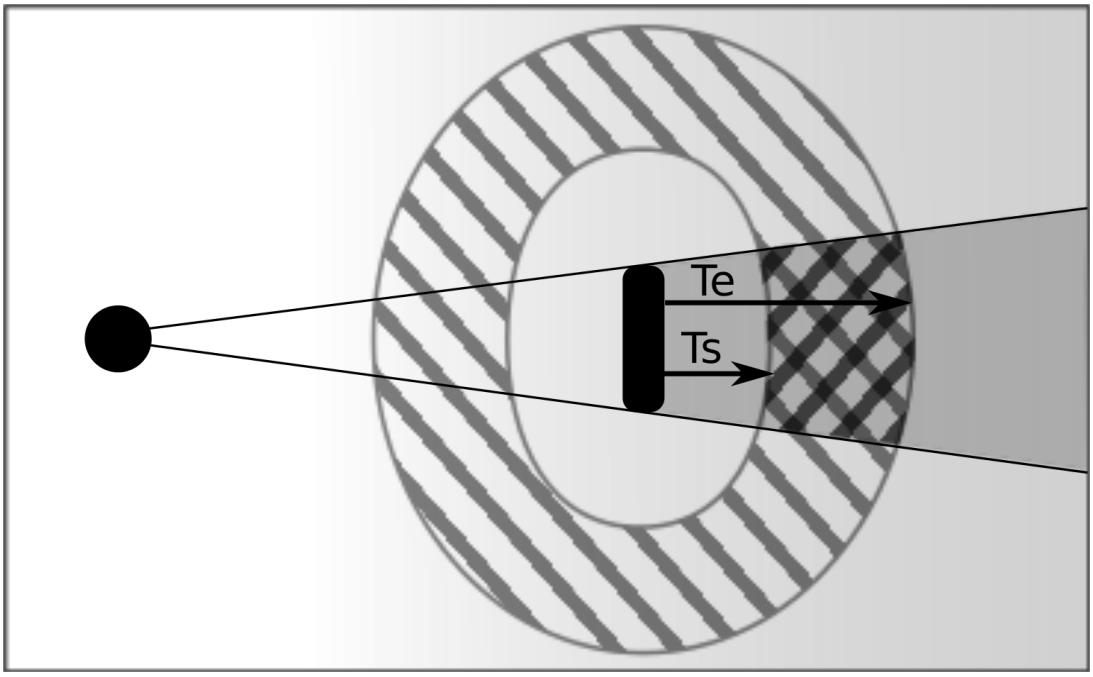

Figure 13.3: Warning spatial structure as intersection of Wave and Shadow

the tag will be reified across the whole Shadow. Most specifically, this is achieved by a new chemical-like rule, matching two annotations located in the same node: the gradient annotation spread from a target p and the annotation of a FE e. The effect of the rule is to update the former annotation, adding e to the list of "tags" it should carry. Additionally, we update the aggregation law such that only Gradient annotations having the same list of tags will be aggregated. In the case of a single FE e, hence, in each node we will possibly have both the annotation tagged by e and the one with empty list of tags, which respectively carry information about a route that crosses an FE and one that circumvents it. In the case of multiple FEs, all combinations will be automatically generated.

Hence, the Shadow structure created from target p due to FE e is defined as the set of nodes whose distance to p along a path crossing e's area is smaller than the one circumventing it, namely, nodes having an annotation tagged with e carrying a smaller distance value that any annotation that is not tagged by e.

Again, while targeting a given point of interest, only agents located in a FE's Shadow should be concerned about the FE, since outside that area an agent would by construction be steered to the target without crossing the FE area. Note that differently from the Horizon Wave pattern, a Gradient Shadow is mostly a static structure, immutable over time: of course, it can still be subject to automatic changes due to mobility of the topology, relocation of gradient source, or interception of new future events.

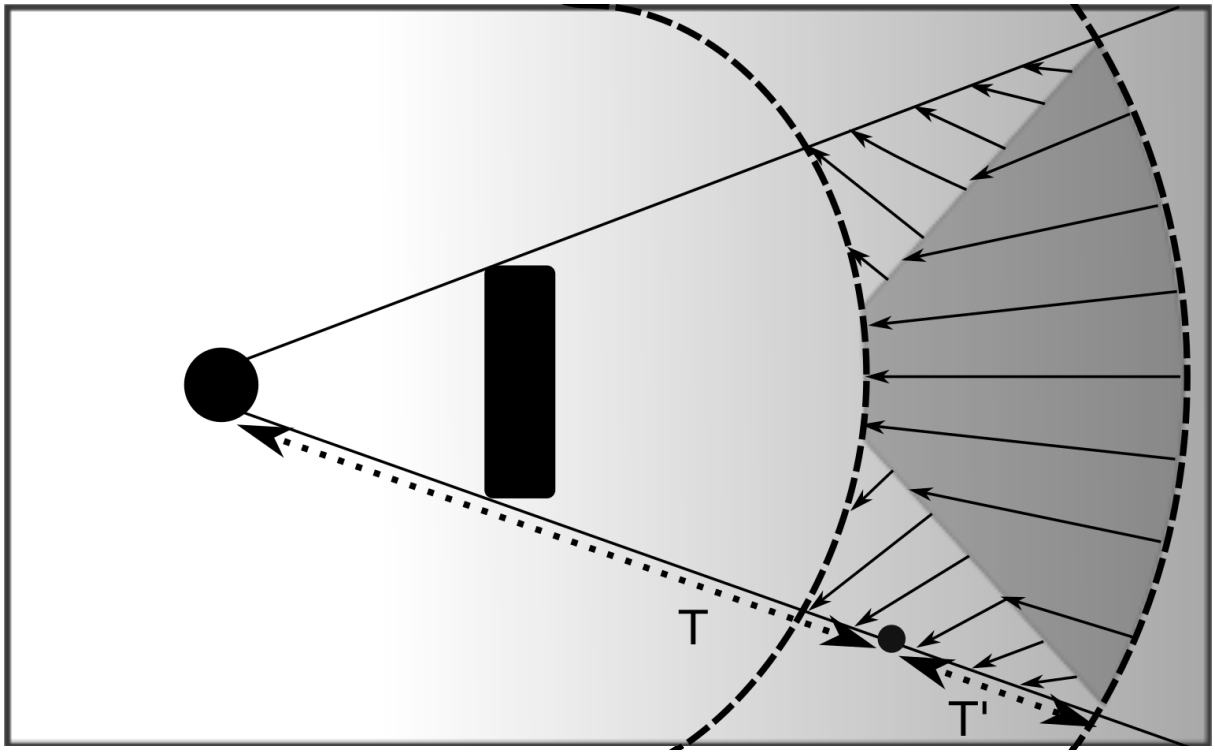

Figure 13.4: Anticipative Gradient pattern (and "waiting distance" T').

Future Event Warning

Now that these two basic bricks have been defined, it is easy to identify their intersection as the portion of space – which we call Future Event Warning (or simply Warning) – in which the Gradient has to be deviated to anticipate the FE. More generally, given a target of interest p, and a set $e_1, .., e_n$ of FEs, let the Global Wave be the union of the Waves of all e_i, then the Warning area is simply obtained by the intersection between p's Shadows as created by $e_1, .., e_n$ and the Global Wave. In our computational framework this area is formed by the nodes featuring the gradient annotation of a target p tagged by an FE e (Shadow), as well as the gradient annotation of FE e whose distance is in between Ts and Te (Wave).

An example of such a structure for a single event is shown in Figure 13.3. Considering only one event e, this area has the shape of a crown sector, it is located on the opposite side of e with respect to the gradient source, and it moves towards e until vanishing.

Anticipative Gradient

We now define the most important self-organisation pattern emerging from the combination of the above ones, namely, what we call the Anticipative Gradient. Starting again from a target p,

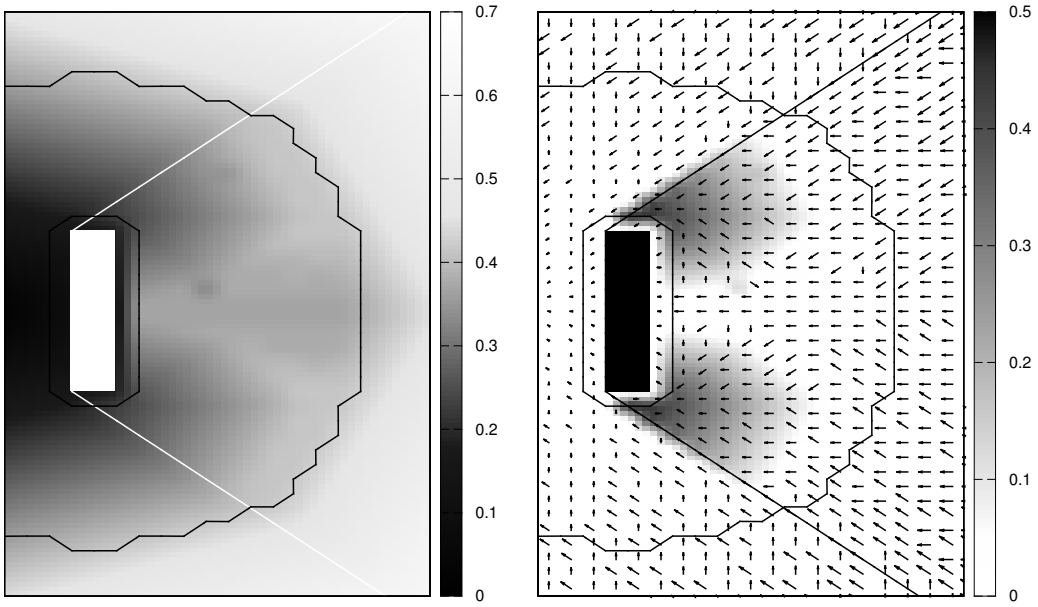

Figure 13.5: An Anticipative Gradient case: (Left) estimated distance normalised by the maximum value; (Right) time-to-destination improvement factor and steering direction.

and a set of future events $e_1, .., e_n$, this is identical to a standard Gradient outside the Warning area, since there agents can be steered to p without any special management. Inside the Warning area of an FE e_i instead, distance and orientation values in each node should reflect the optimal path among the two following alternative: the one crossing e_i and the one circumventing it.

An additional caveat is how we can compute the actual time-to-destination (i.e., distance) of a node n in the Warning area along a path crossing an FE e—since using that path will make the agent experience the obstacle in the future. A possibility for an agent would be to wait in n until it is outside the Warning area and then move towards e and then to p. Alternatively, it could immediately move from n towards e, and be stuck there until the Wave is over—it can be shown that the strategy of circumventing the FE area instead of staying stuck is never winning over the one finding a path circumventing e since the beginning. In any strategy, the time-to-destination is the one obtained without considering e (T), plus a *waiting distance* (T') measuring the distance of n to the outside margin of the Wave (namely, the waiting time the agent would be forced to)—see Figure 13.4. Note that in the case n is simultaneously inside many different Warning areas, the distance value should be increased by the waiting distance computed for each of them.

Figure 13.4 also qualitatively depicts what are expected directions in a Warning area. On the one hand, near the end of the Wave and in central position there is relatively little advantage in trying to circumvent the future event location, for the strategy of moving towards the FE anyway is winning. Hence, the orientation is towards the FE. In that region, we expect a rather uniform

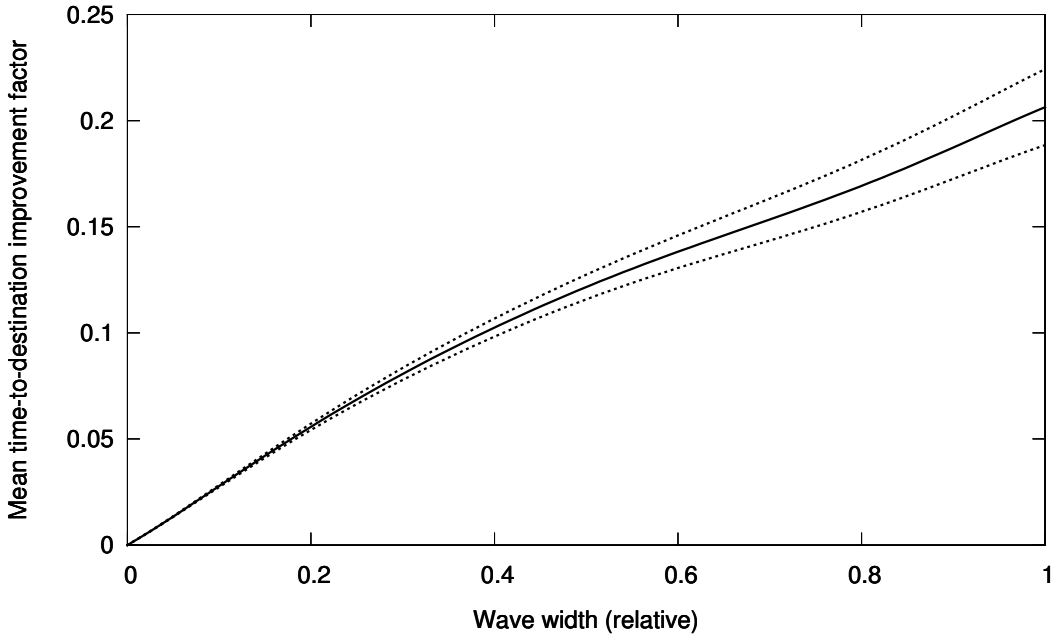

Figure 13.6: Mean time-to-destination factor (and standard deviation) in the Warning area, measured with different widths of Horizon Wave (normalised to grid size).

time-to-destination. On the other hand, at the beginning of the wave and in lateral positions of the Warning area the steering service will likely redirect an agent laterally outside the Warning zone to bypass the FE area.

More precisely, the Anticipative Gradient can be created out of two additional laws working as follows. The first law simply takes the Gradient annotation generated from a target, and clones it into a unique new annotation for the Anticipative Gradient: at this stage the two annotations will report the same distance D from the source. The second law matches the annotation a of an Anticipative Gradient as created above with the annotation w of the Wave of an FE e: its result is that e is removed from the list of tags in a, and a's distance is increased by e's waiting distance as discussed above, computed as $T_e - D$. Note such a law would iteratively fire to match a with the waves of all active (overlapping) FEs, until a's list of tags is emptied.

The neat result is that the Antipative Gradient is identical to the standard Gradient outside Warning areas, it instead selectively reports augmented distances in Warning areas, such that following the route towards nodes with minimum distance to the target always guarantee to efficiently avoid obstacles appearing in the future.

13.1.3 Simulations

We checked the basic correctness of the proposed self-organisation pattern by simulation, conducted using ALCHEMIST.

Figure 13.5 (left) shows a pictorial representation of how the Anticipative Gradient looks like in a simulation of a 2D grid of 60x30 nodes, with a white-coloured rectangular-like FE. Figure 13.5 (right) shows what would be the right direction proposed at that time in each node. The analysis of these data confirms the expected qualitative results, in analogy with Figure 13.4. Namely, in the Warning area there is a whole triangle-like sector in which the estimated distance is higher (lighter colour) and constant—due to the waiting-time an agent has to experience. There, the proposed direction does not circumvent the FE. Instead, in the upper and lower border of the Warning area the Anticipative Gradient would tend to steer an agent away from the FE. Figure 13.5 (right) also shows what is the relative improvement an agent would experience when starting from the different locations of the network thanks to the anticipation of the future event. It ranges from a 50% in nodes near the left border of the Warning area to a 0% in nodes outside the Warning area.

Figure 13.6 shows what is the average (with standard deviation) improvement for nodes in the Warning area depending on the width of the Wave. Clearly, there is a greater improvement in using the Anticipative Gradient pattern in those cases where the future event lasts longer, namely, with higher width ($Te - Ts$) of the Wave.

14 Case Study: Crowd evacuation

14.1 A crowd evacuation application

A public exposition is taking place in a museum composed of rooms, corridors, and exit doors. Figure 14.1 shows the layout of the exposition. The surface of the exposition is covered by sensors, arranged in a grid, able to sense fire, detect the presence of people, interact with other sensors in their proximity as well as with PDAs that visitors carry with them.

When a fire breaks out, PDAs (by interaction with sensors) must show the direction towards an exit, along a safe path. The system has to be resilient to changes or unpredicted situations, in particular the safe path must consider: *(i - distance)*: it should tend to lead to the nearest exit; *(ii - fire)*: it should tend to stay away from fire; and *(iii - crowd)*: it should tend to avoid overcrowded paths. These factors influence PDAs by means of the following LSAs:

- The exit gradient: each node contains an LSA with a numeric value indicating the distance from the nearest exit. These LSAs form a gradient field covering the whole expo. Over an exit, the gradient value is 0, and it increments with distance from the exit. When there are several exits or several paths towards an exit, the lowest distance value is kept for a node.
- The fire gradient: each node also contains an LSA indicating its distance from nearest fire. The LSA's value is 0 at the fire's location and increases with distance. It reaches a

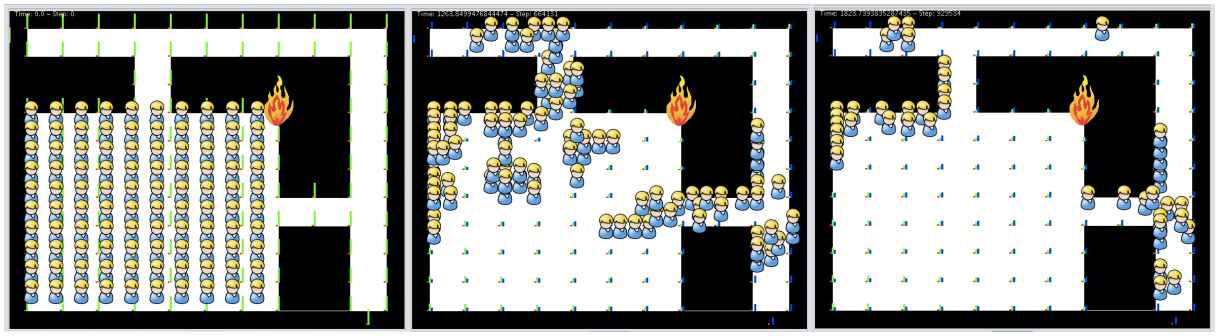

Figure 14.1: A simulation run of the reference exposition: three snapshots

maximum value at a safe distance from a fire.

- The crowding gradient: sensors in the expo surface detect the presence of people and adjust the value in a local LSA indicating the crowding level of the location. As with exit and fire, this LSA is diffused around, and its value increases with the distance from the crowded region.

- The attractiveness information: finally, each node contains an LSA indicating how desirable it is as a position in an escape route. Its value is based on the values of the previous three, and is used to choose which direction displays should point to.

14.1.1 Types of LSAs in the system

There are three forms of LSAs used in this scenario:

$$\langle \texttt{source}, type, max, ann \rangle,\ \langle \texttt{grad}, type, value, max, ann \rangle,\ \langle \texttt{info}, type, value, tstamp \rangle$$

A **source** LSA is used for gradient sources: *type* indicates the type of gradient (`fire`, `exit`, and so on); *max* is the gradient's maximum value; and *ann* is the annealing factor—its purpose will be described later, along with eco-laws. A **gradient** LSA is used for individual values in a gradient: *value* indicates the individual value; and the other parameters are like in the source LSAs. Finally, an **info** LSA is used for local values (e.g., not part of a gradient)—parameters are like in the source and gradient LSAs. The *tstamp* reflects the time of creation of the LSA.

14.1.2 Building the fire, exit and crowding gradients

The sources of the gradients are injected by sensors at appropriate locations, with the values $\langle \texttt{source}, \texttt{exit}, \texttt{Me}, \texttt{Ae} \rangle$ and $\langle \texttt{source}, \texttt{fire}, \texttt{Mf}, \texttt{Af} \rangle$. For the crowding information, we may assume that sensors are calibrated so as to locally inject an LSA indicating the level of crowding. The actual threshold in number of people will mainly depend on the sensor arrangement, and should be seen as a configuration issue. The crowding LSA will look like $\langle \texttt{source}, \texttt{crowd}, \texttt{Mc}, \texttt{Ac} \rangle$ and is periodically updated by the sensor.

As sources are established, gradients are built by the following set of eco-laws, applying to exit, fire and crowding gradients by parameterising the *type* in the LSAs. First, we define eco-law 14.1 that, given a source, initiates its gradient:

$$\langle \texttt{source}, T, M, A \rangle \overset{R_{init}}{\longmapsto} \langle \texttt{source}, T, M, A \rangle, \langle \texttt{grad}, T, 0, M, A \rangle \tag{14.1}$$

When a node contains a gradient LSA, it spreads it to a neighbouring node with an increased value, according to eco-law 14.2:

$$\langle \texttt{grad}, T, V, M, A \rangle \overset{R_s}{\longmapsto} \langle \texttt{grad}, T, V, M, A \rangle, +\langle \texttt{grad}, T, \texttt{min}(V+\#D, M), M, A \rangle \tag{14.2}$$

Due to use of the #*D* system variable, each node will carry a `grad` LSA indicating the topological distance from the source. When the spread values reach the maximum vale M, the gradient becomes a plateau. Also note that the iterative application of this eco-law causes continuous diffusion of the LSA to all neighbouring nodes.

The spreading eco-law above may produce duplicate values in locations (due to multiple sources, multiple paths to a source, or even diffusion of multiple LSAs over time). Thus, eco-law 14.3 retains only the minimum distance:

$$\langle \texttt{grad}, T, V, M, A\rangle, \langle \texttt{grad}, T, W, M, A\rangle \rightarrow \langle \texttt{grad}, T, \texttt{min}(V,W), M, A\rangle \tag{14.3}$$

Finally, we have to address the dynamism of the scenario where people move, fires extinguish, exits may be blocked, crowds form and dissolve. If a gradient source vanishes, the diffused values should increase (e.g., the distance to exit increases if the nearest exit is no longer available). However, with the above eco-laws this does not happen because eco-law 14.3 always retaining the minimum value. This is the purpose of the annealing parameter in the gradient LSAs: it defines the rate of eco-law 14.4, which continuously tends to level up gradient values, encouraging the replacement of old values by more current ones:

$$\langle \texttt{grad}, T, V, M, A\rangle \xmapsto{R_{ann}(A)} \langle \texttt{grad}, T, V+1, M, A\rangle \tag{14.4}$$

The R_{ann} rate is directly proportional to A. When a fire is put out, for example, this eco-law will gradually raise the fire gradient to the point where it reaches the maximum, indicating no fire. Annealing may introduce a burden on the system, therefore high annealing values should only be used for gradients that have to change often or quickly.

14.1.3 Ranking escape paths: the attractiveness field

Based on exit distance, fire distance and crowding, a location can be ranked as more or less "attractive" to be part of an escape path. This is done via an *attractiveness* value automatically attached to each node by eco-law 14.5:

$$\begin{aligned}&\langle \texttt{grad}, \texttt{exit}, E, Me, Ae\rangle, \langle \texttt{grad}, \texttt{fire}, F, Mf, Af\rangle, \langle \texttt{info}, \texttt{crowd}, CR, TS\rangle \xmapsto{R_{att}} \\ &\langle \texttt{grad}, \texttt{exit}, E, Me, Ae\rangle, \langle \texttt{grad}, \texttt{fire}, F, Mf, Af\rangle, \langle \texttt{info}, \texttt{crowd}, CR, TS\rangle, \\ &\langle \texttt{info}, \texttt{attr}, (Me-E)/(1+(Mf-F)+k\times(Mc-C)), \#T\rangle\end{aligned} \tag{14.5}$$

Coefficient k (tuned by simulation) is used to weight the effect on crowding on attractiveness. As gradients evolve, older attractiveness LSAs are replaced with newer ones (T is assumed positive):

$$\langle \texttt{info}, \texttt{attr}, A, TS\rangle, \langle \texttt{info}, \texttt{attr}, A2, TS{+}T\rangle \rightarrow \langle \texttt{info}, \texttt{attr}, A2, TS{+}T\rangle \tag{14.6}$$

14.1.4 Choosing a direction

Each location contains by default an LSA of the form $\langle \texttt{info}, \texttt{escape}, L, TS \rangle$, where L is the direction to be suggested by the PDA. In principle, the neighbour with the highest attractiveness should be chosen, but a more resilient solution is to tie the markovian rate of eco-laws to the attractiveness of neighbours, so that the highest probability is to point the best neighbour, with a possibility to point a less-than-optimal (but still attractive) neighbour:

$$\begin{array}{l} \langle \texttt{info}, \texttt{escape}, L \rangle, \langle \texttt{info}, \texttt{attr}, A, TS \rangle, +\langle \texttt{info}, \texttt{attr}, A+\Delta, TS2 \rangle \xmapsto{R_{disp}(\Delta)} \\ \langle \texttt{info}, \texttt{escape}, \#O \rangle, \langle \texttt{info}, \texttt{attr}, A, TS \rangle, +\langle \texttt{info}, \texttt{attr}, A+\Delta, TS2 \rangle \end{array} \tag{14.7}$$

The rate is proportional to the difference in attractiveness between the node and its neighbour (Δ). The higher the Δ, the higher the rate. Note that Δ is a positive value, hence $A+\Delta$ implies that eco-law 14.7 only considers neighbours with a higher attractiveness, i.e., the PDA will not point *away* from the exit.

14.1.5 Resilient behaviour

The proposed architecture is intrinsically able to dynamically adapt to unexpected events (like node failures, network isolation, exits suddenly unavailable, crowd formation, and so on) while maintaining its functionality. We will now discuss a few examples of possible problems and explain how the system reacts to them.

Single-node failure

Most behaviour in this architecture is based on nodes being able to apply eco-laws and host LSAs. Failure of a node clearly impacts this behaviour in some measure. A failing node (i.e., disappearing from the system) results in a physical location where no information for the displays is available, and where gradient information is not received or transmitted.

If the failing node is a generic one (i.e, no fire or exits at that node) and its disappearance does not cause network isolation, the impact is somewhat limited. Gradients will still be transmitted *around* the broken node, although values will change. Displays traversing a failing node will not update their direction; however, they will also never guide people towards a failing node, because it does not have an attractiveness value; they will rather steer people around the failing node. This means that functionality is preserved with decreased efficiency (i.e., longer paths).

If the failing node is a node containing a fire, the consequences are more severe: the fire will not be detected. PDAs will still never point the failing node on fire directly, but they will let people pass near it. One might argue that people will still see the fire; however, from the point of view of system behaviour, safety is reduced. On the other hand, it can be assumed that a serious fire would be detected by more sensors, and that all of them failing at the same time is unlikely.

If the failing node is an exit node, functionality gets a major hit, because the whole system will lack a fundamental information for display functioning (the exit gradient). If other exits

are available, the system will still guide people towards them (again preserving functionality with reduced efficiency). However, if the failing node "hides" the only available exit, the system completely fails to provide useful information. This situation can be tackled by redundant source nodes near each exit.

Network isolation

When a group of failing nodes is the only connection between two parts of the network, this causes network isolation. There are a few sub-scenarios to consider. If the isolated section does not include exits, we find another limitation: the evacuation mechanism cannot work inside that section, as it is lacking fundamental information. If the section does however contain exits, it will act as an autonomous system in itself. Efficiency may be reduced but functionality will still be preserved, with an important caveat: this section of the network will ignore fires and alarms occurring in the rest of the network. This may or may not be an acceptable limitation depending on the scenario. For fires, it could probably be unacceptable. In a general way, however, we would consider that an isolated part of the network behaving as an autonomous system is an acceptable fallback. Again, redundancy is desirable for those nodes of the network that can cause isolation.

Exit unavailability

If an exit is suddenly unavailable in the expo, for instance it is broken and hence it does not open, a sensor could detect this and drop the source LSA. Because of the annealing mechanism, this causes the gradient value in each node to converge to the situation corresponding to the unavailability of that door; namely, a new path toward another exit emerges and affects the behaviour of all PDAs—people would simply go back to another exit as expected.

Crowd formation

While people follow their PDAs towards the nearest exit, it is possible that a jam forms in front of certain doors or corridors: this is one of the most perilous, unexpectable situations to appear in crowd evacuation. The application we set up is meant to emergently tackle these situations by means of the crowding gradient.

Assume two available corridors exist towards some exit, and from a given room people can choose to walk through any of them towards the exit. If one becomes unavailable because a jam is forming the crowding gradient would emerge, reducing attractiveness along that corridor. Accordingly, it is expected that people behind the crowded area start walking back and head for the other corridor, and people in the room start choosing systematically the path free from crowd, even if they do not "see" the problem from the room.

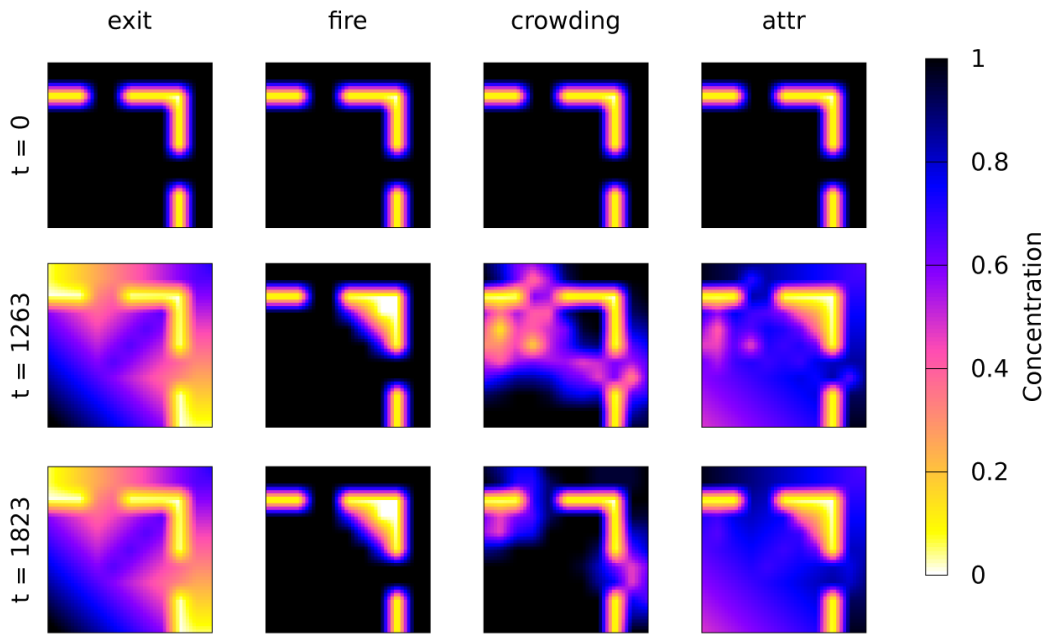

Figure 14.2: Gradients for the snaphots in Figure 14.1. Concentration is the gradient value normalised to its maximum value.

14.2 Simulation in ALCHEMIST

The proposed system is specified by a limited set of eco-laws, equipped with CTMC semantics. This specification can be used to feed the SAPERE infrastructure and be simulated in order to check the model dynamics in advance.

14.2.1 Simulation setting

We here present simulations conducted over an exposition structured as shown in Figure 14.1, where three snapshots of a simulation run are reported: all the people in the room start moving towards one of the two exists (located at the ends of the corridor) because of the fire in the top-right corner of the room. Note in third snapshot that a person is walking in the middle of the corridor, for she was suggested to go to a farther exit because of the corridor jam at the bottom-right.

Rooms and corridors are covered by a grid of locations hosting sensors, one per meter in the room, one per two meters in the corridor: such locations are the infrastructure nodes where LSAs are reified. The maximum values for the gradients are set to: $\texttt{Me} = 30$, $\texttt{Mf} = 3$, $\texttt{Mc} = 20$. The PDA of each person is modelled as a mobile node, able to perceive the attractiveness gradient in the nearest sensor location: accordingly, the person moves in the suggested direction.

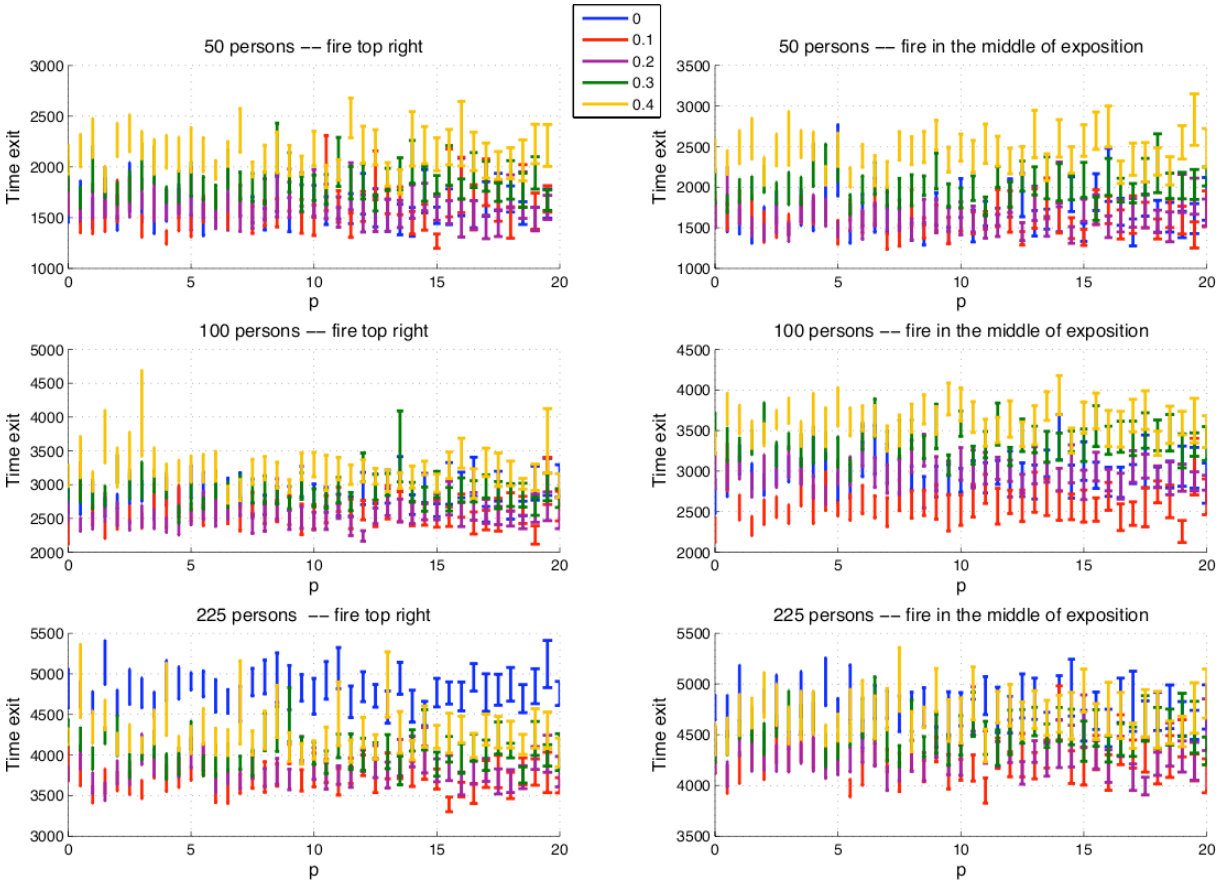

Figure 14.3: Results of parameter analysis

Figure 14.2 shows the gradients of exit, fire, crowding and attractiveness (one per column) corresponding to the simulation steps of Figure 14.1 (one step per row). At $t = 0$, gradients are level; with time, the gradient self-modify—it is easy to see the exits, fire and crowds in the respective gradients. The crowding gradient in the third column changes dynamically during simulation according to the movement of people. The last column shows the attractiveness gradient, computed from the other three gradients. Note how the second snapshot shows an attractiveness "hole" in the middle of the room and in the corridor due to crowding.

14.2.2 Tuning parameters

We chose to tune by simulation two parameters of the model: k (used in the eco-law computing attractiveness), and a multiplication factor p we applied to the crowding gradient to control its slope—namely, how far crowding should be perceived. Different simulations have been performed in order to find the set of parameters that minimises the time of exit of all the people in

the room: the range of variation of k is [0.1, 1] with a step of 0.1, while the range for p is [0, 20] with step of 0.5. The smaller is k, the lower is the impact of the crowding in the computation of attractiveness, so that we expect the crowds not be avoided and the exit time to grow. The bigger is p, the higher is the crowding gradient slope, until it becomes a local information that PDAs can perceive only very close to it. Experiments have been done with 50, 100 and 225 persons in the room and with a fire in the corner of the room or in the middle, so as to experimentally observe the impact of parameters in the evacuation outcome of different scenarios. For each couple of parameters, 10 simulations have been run, considering average value and standard deviation of the resulting time.

In Figure 14.3 we show the results we obtained from the analysis of the parameters. Only dynamics for $k < 0.5$ are shown, because the time of exit is much higher elsewhere, and also because standard deviation becomes much higher since the system is much more chaotic.

The graphs show that with $k = 0$ (i.e., if the crowding gradient does *not* influence attractiveness) the performance of the system slows down as the number of persons increases, and the likelihood of forming of crowds grows. This result highlights the impact of considering crowds in resilience with respect to jams. For $k = 0.1$ we obtained the best performance for the system for both simulated fire positions. Results finally showed that the parameter p does not seem to have relevant impact on the average exit time in this particular expo.

14.2.3 Resilience to node failures

Among the many simulations analysing resilience aspects, we focused on the case in which all nodes experience temporary failures, i.e., they work correctly only for a percentage of time. We expect the system to still be safe, i.e., be able to drive people out of the room in reasonable time, in spite of the changes in gradient values around the broken nodes. The results of simulation with 50, 100 and 225 persons are shown in Figure 14.4: the time to exit does not significantly increase even in the case of nodes failing 30%-40% of the time, and people are still able to eventually reach the exit even with 80% of node downtime.

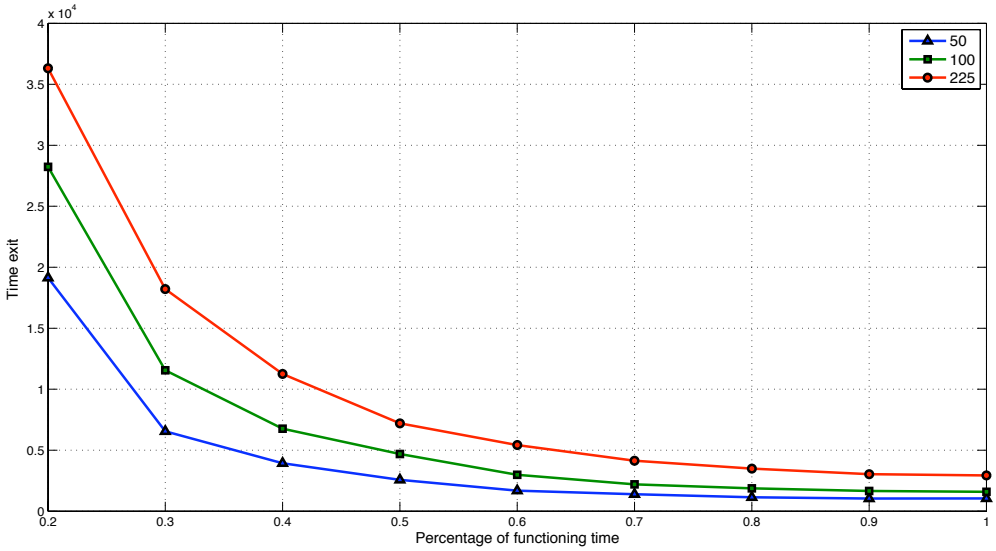

Figure 14.4: Resilience to temporary node failures. Numbers in the key represent the number of people.

15 Case Study: Crowd steering

In this chapter we show a crowd steering case study in which a middleware has the goal of leading people in the desired location within a complex environment in short time, avoiding obstacles such as crowded regions and without global supervisioning.

15.1 Reference scenario

Consider a museum with a set of rooms, whose floor is covered with a network of computational devices (infrastructure nodes). These devices can exchange information with each other based on proximity, sense the presence of visitors, and hold information about expositions currently active in the museum. Each room has four exits and they are connected via external corridors. Visitors wandering the museum are equipped with a hand-held device that holds the visitor's preferences. By interaction with infrastructure nodes, a visitor can be guided towards rooms with a target matching their interest, thanks to signs dynamically appearing on his smartphone or on public displays. This is done using techniques suggested in the field of spatial computing [VCMZ11]—namely, computational gradients injected in a source and diffusing around such that each node holds the minimum distance value from the source.

The environment is a continuous bidimensional space with walls. Smartphones (or public displays) are agents dynamically linked with the nearest infrastructure node – the neighbours are the sensors inside a certain radius r, parameter of the model – from which they can retrieve data in order to suggest visitors where to go. Visitors are agents which follow the advices of their hand-held device (or public displays). It is defined a minimum possible distance between them, so as to model the physical limit and the fact that two visitors can't be in the same place at the same time. Visitors can move of fixed size steps inside the environment. If an obstacle is on their path, their movement is shortened to the allowed position nearest to the desired place.

15.2 Steering strategy

All the information exchanged is in form of annotations, simply modelled as tuples $\langle v_1, \ldots, v_n \rangle$ (ordered sequence) of typed values, which could be for example numbers, strings, structured

types, or function names. There are three forms of annotations used in this scenario:

$$\langle \texttt{source}, id, type, N_{max} \rangle$$
$$\langle \texttt{field}, id, type, value, tstamp \rangle$$
$$\langle \texttt{info}, id, crowd, M, t' \rangle$$

A **source** annotation is used as a source with the goal of generating a field: *id* labels the source so as to distinguish sources of the same type; *type* indicates the type of fields in order to distinguish different expositions; N_{max} is the field's maximum value. A **field** annotation is used for individual values in a gradient: *value* indicates the individual value; the *tstamp* reflects the time of creation of the annotation; the other parameters are like in the source annotation. An **info** annotation is supposed to be created and kept up to date by each sensor. M represents the number of smartphones the sensor is perceiving as neighbours.

The rules are expressed in form of chemical-resembling laws, working over patterns of annotations. One such pattern P is basically an annotation which may have some variable in place of one or more arguments of a tuple, and an annotation L is matched to the pattern P if there exists a substitution of variables which applied to P gives L. A law is hence of the kind $P_1, \ldots, P_n \overset{r}{\mapsto} P'_1, \ldots, P'_m$, where: *(i)* the left-hand side (reagents) specifies patterns that should match annotations $L_1, \ldots, L_n$ to be extracted from the local annotation space; *(ii)* the right-hand side (products) specifies patterns of annotations which are accordingly to be inserted back in the space (after applying substitutions found when extracting reagents, as in standard logic-based rule approaches); and *(iii)* rate r is a numerical positive value indicating the average frequency at which the law is to be fired—namely, we model execution of the law as a CTMC transition with Markovian rate (average frequency) r. If no rate is given the reaction is meant to be executed "as soon as possible", which means that the rate that associated with the reaction tends to infinite. To allow interaction between different nodes (hence, annotation spaces), we introduce the concept of *remote pattern*, written $+P$, which is a pattern that will be matched with an annotation occurring in a neighbouring space.

As sources annotations are injected in nodes, gradients are built by the first three rules in Figure 15.1. The first one, given a source, initiates the field with its possible maximum value. The second one, when a node contains a field annotation, spreads a copy of it to a neighbouring node picked up randomly with a new value computed considering the crowding around the sensor. The parameter k allows to tune how much crowding should influence the field, while $\#D$ is the measure of the distance between the two involved sensors. As a consequence of these laws, each node will carry a field annotation indicating the topological distance from the source. The closest is the field value to N_{max}, the nearest is the field source. When the spread values reach the minimum value 0, the gradient has to become a plateau.

To address the dynamism of the scenario where people move, targets being possibly shifted, and crowds forming and dissolving, we introduced the following mechanism. We expect that if a gradient source moves the diffused value has to change according to the new position. This is the purpose of the *tstamp* parameter which is used in the fourth law, continuously updating old

$$\langle \texttt{source}, id, type, N_{max} \rangle \xmapsto{r_{init}} \langle \texttt{source}, id, type, N_{max} \rangle, \langle \texttt{field}, id, type, N_{max}, \#T \rangle$$

$$\langle \texttt{field}, id, type, N, t \rangle \xmapsto{r_{diff}} \langle \texttt{field}, id, type, N, t \rangle, +\langle \texttt{pre\_field}, id, type, N - \#D, t \rangle$$

$$\langle \texttt{pre\_field}, id, type, N, t \rangle, \langle \texttt{info}, id, crowd, M, t' \rangle \mapsto \langle \texttt{field}, id, type, N - k * C, t \rangle, \langle \texttt{info}, id, crowd, M, t' \rangle$$

$$\langle \texttt{field}, id, type, N, t \rangle, \langle \texttt{field}, id, type, M, t' + t \rangle \mapsto \langle \texttt{field}, id, type, M, t' + t \rangle$$

$$\langle \texttt{field}, id, type, N, t \rangle, \langle \texttt{field}, id, type, M, t \rangle \mapsto \langle \texttt{field}, id, type, max(M, N), t \rangle$$

$$\langle \texttt{field}, id, type, N, t \rangle, \langle \texttt{field}, id', type, N + M, t' \rangle \mapsto \langle \texttt{field}, id', type, N + M, t' \rangle$$

Figure 15.1: Laws describing the museum application.

values by more recent ones (*youngest* law). In this way we ensure that the system is able to adapt to changes of the source states. Finally, the spreading law above will produce duplicate values in locations, due to multiple sources of the same type (indicated by different ids), multiple paths to a source, or even diffusion of multiple annotations over time. For this reason we introduced the last two laws. They retain only the maximum value, i.e. the minimum distance, the former when there are two identical annotations with only a different value, the latter when the id is different (*shortest* laws). The proposed solution is intrinsically able to dynamically adapt to unexpected events (like node failures, network isolation, crowd formation, and so on) while maintaining its functionality.

People are modelled as nodes programmed with a single reaction with no conditions and having, as action, the ability to follow the highest field value among neighbouring sensors. For the crowding annotation, we may assume that sensors are calibrated so as to locally inject and periodically update it by setting the current level of crowding, *i.e.* the number of persons. The reaction rates are identified by hand performing different simulations with different parameters. The results reported are obtained with $r_{init} = 1$ and $r_{diff} = 50$. The other laws show no rate because it is assumed to be infinite (as-soon-as-possible semantics).

15.3 Simulation in ALCHEMIST

We here present simulations conducted over an exposition, where nine rooms are connected via corridors. People can express different preferences represented by their shape.

Three snapshots of a first simulation run are reported in Figure 15.2. We here consider four different targets that are located in the four rooms near environment angles. People are initially

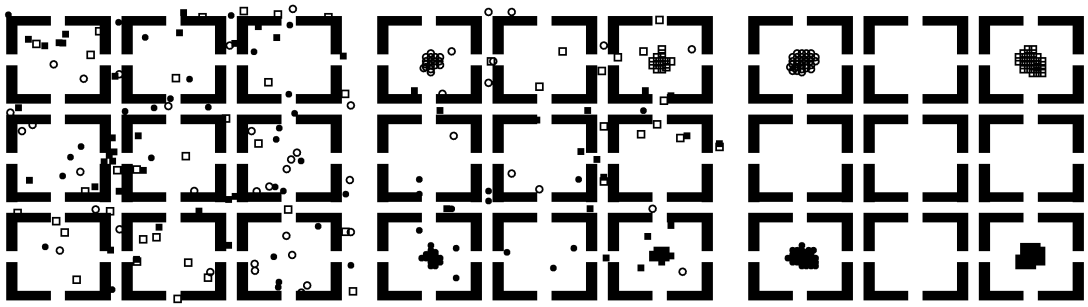

Figure 15.2: A simulation run of the reference exposition: three snapshots of ALCHEMIST's graphic reporting module with this simulation.

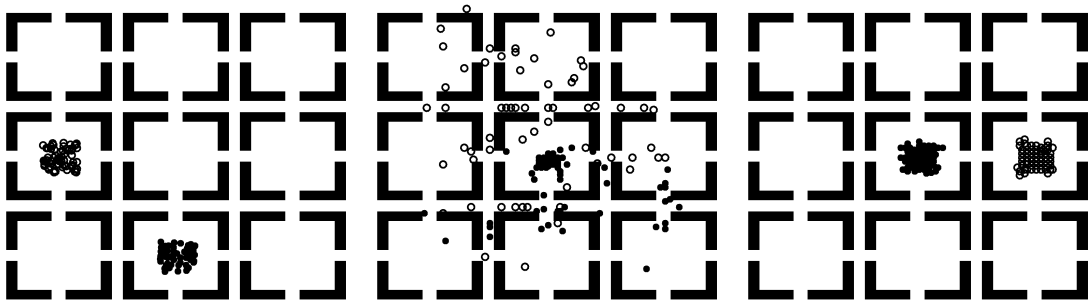

Figure 15.3: A run showing the effect of crowding: dark visitors occupy a central room, making other visitors moving left to right by a longer, less crowded path

spread randomly in the museum, as shown in the first snapshot, and they eventually reach the room in which the desired target is hosted, as shown in the last snapshot.

Figure 15.3 shows a simulation experimenting with the effect of crowding in the movement of people. Two groups of people – denoted with empty and full circles – with common interests are initially located in two different rooms, as shown in the first snapshot. The target for the dark visitors is located in the central room of the second row, while the others' is in the right room of the second row. In the simulation, dark visitors reach their target soon because it is nearer, thus forming a crowded area intersecting the shortest path towards the target for the other visitors. Due to this jam the latter visitors choose a different path that is longer but less crowded.

Both tests show qualitative effectiveness of the proposed laws, and suggest that our simulation approach can be used for additional experiments focussing on tuning system parameters (factor k) or alternative strategies (e.g., diffusing crowd information) to optimise paths to destinations. For instance, in the context of the second case, Figure 15.4 shows how factor k can influence the time for (sub)groups of (light) people to reach the destination, by which we can see that even small values of k lead to a significant improvement—which slowly decreases as k grows.

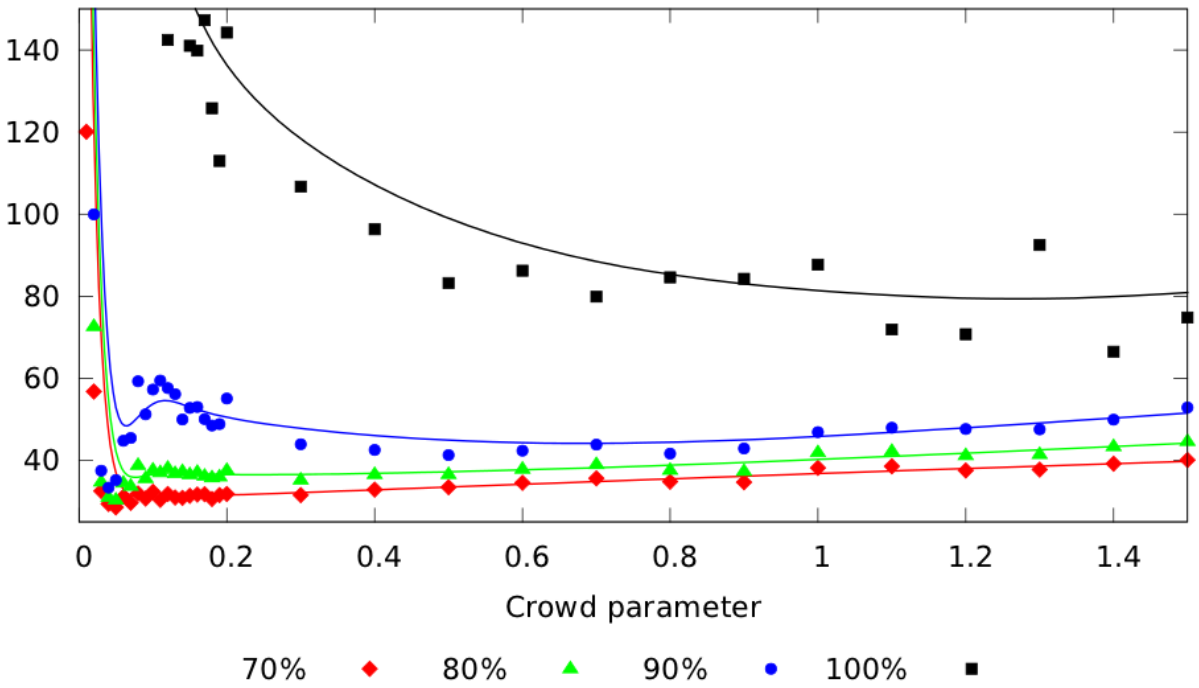

Figure 15.4: Time units of convergence time with different values of crowd parameter. For each line, we measure the time required for a certain percentage of people to arrive to destination. For instance, the red line measures the time required for 70% of the group to arrive to destination.

16 Case study: Self composition

Services are one of the main actors in pervasive and ubiquitous computing. Composing services is a big issue in these contexts, providing new applications or high level services. Self-composition is even more challenging.

In literature some surveys and possible solutions have been proposed, but only few of them give concrete proposals for the design, implementation and evaluation of service composition. Moreover the composition of services is solved in a centralised way, solution which hardly deal with the large-scale distributed property of pervasive systems, or the composition is specified by the application developers, preventing the system to use new services that appear at runtime.

In this chapter we frame/situate the problem of self-composition in the framework of pervasive ecosystems, defining what a service is and what composing services means. We then propose an approach to perform self-composition of services and demonstrate our ideas showing how gradients can be composed with other services. In particular our discussion focuses on the composition of gradient with services providing local levels of crowd, at the purpose of avoiding path crossing overcrowded areas. This example is modelled in terms of LSAs and eco-laws, that are the basic components of a pervasive ecosystem [VPMS12], and preliminary simulations, aimed at validating the approach, are run on top of ALCHEMIST.

16.1 Self-Composition approaches

A software service represents a functionality that can be provided on demand in order to create higher level software artefacts, such as applications or higher level services. The notion of service allows applications to be easily developed and to adapt by changing the different services, which the application is composed by, at run-time. Self-composition of services involves the automatic discovery of new services by composing available ones, the selection of services, the plan to execute them (i.e. services execution order), and the adaptation of the composed service when new requirements appear or a service disappears from the system.

Service Composition in SOA

The static character of traditional composition approaches such as orchestration and choreography has been recently challenged by so-called dynamic service composition approaches, involving semantic relations [BBG06], and/or AI planning techniques to generate process automatically based on the specification of a problem [WRG+07]. Their basic goal is to analyse a description of the services to compose and compute a composition satisfying structural/behaviour/ontological compliance of the result of composition as can be deduced from information about the composites. One of the main challenges of these approaches is their limited scalability and the strong requirements they pose on the detail of service description.

Evolutionary techniques

Evolutionary approaches such as those based on Genetics Algorithms (GA) have been proposed as well for service composition [CDPEV05], motivated by the need of determining the services participating in a composition that satisfies certain QoS constraints, which is known to be a NP-Hard problem. More generally, evolutionary approaches have been used in the field of autonomic computing to establish the set of norms, policies or rules that drive the system to the desired emergent behaviour [SRAA10]. We observe that this approach has a potential for service composition: in the same way as the genetic programming determine the proper execution of functions and their parameters, it could be used to find which services, their order of execution and the parameters that its service should receive.

Competition-based approaches

In order to tackle the potentially high number of different compositions that can arise in an open system, a mechanism of *coordination reactions* for networked tuple spaces is proposed [Vir11], showing how it can be applied to support competition and composition in a fully distributed setting—services opportunistically compose in those regions of the network where this results in a more competitive service. While all compositions can possibly take place, thanks to a positive/negative feedback only successful ones are meant to "survive" in the system, the others fading until becoming never actually selected. This is achieved by matching services with requests through a self-regulating dynamics inspired by prey-predator model of population dynamics [Ber92b], and by diffusing service tuples using a computational-field approach [MZ05, BB06].

16.2 A comprehensive approach for pervasive ecosystems

As previously emphasised, the issue of service composition is becoming wider and wider, and is now encompassing new research areas. Among them – pervasive computing, autonomic computing, robotics – we here focus on pervasive computing, and in particular, on the framework

of pervasive ecosystems [ZV11]. There, we understand and design a very large and dense pervasive computing system as a sort of substrate in which humans, institutions, software systems and devices, inject services of various kinds without an a-priori knowledge of the structure and behaviour of those already available and those that will be injected later. Such services diffuse in the network, and are situated and context-aware in the sense that they contextualise to (i.e., they will affect and be affected by) the situation in each *niche* of it. In turn, while in standard SOA a service is a centralised point of functionality, interacting with its clients by stateful protocols of message-passing actions, a service is here a more generalised concept: it is any activity, triggered by the intention of some localised software agent, that flows in the system updating the spatial and temporal configuration of other services, and generating a set of events ultimately perceived by the other agents in the system. Examples of such services include: the materialisation of data produced by all sensors available around, routing services able to interconnect devices based on their physic al, logical or social proximity, local/global advertisers of the availability of some content, situation recommenders capturing relevant information about a context, and so on. Typical application scenarios include pervasive display ecosystems, augmented social reality, and traffic control—the reader is deferred to [ZV11] for a more deeper description of them.

16.3 The self-composition issue in pervasive service ecosystems

Pervasive ecosystems naturally call for a radical shift in the way software is designed and developed. The concept of "software system" loses its standard meaning of a monolitically designed/implemented/deployed artefact. It is rather a mash-up of services, and the development of a "new software system" simply translates into the development of additional services to be immersed in a scenario of existing ones, with the goal of composing with them and evolving their functionalities, thus making software live in a sort of "eternal beta" state [KC09].

The natural questions we have to answer in order to support this vision are hence: what can make services compose if they were not explicitly designed to do so? How can we decide "whether" and "how" two or more services should compose? How can such decisions be continuously reconsidered depending on the actual (spatial/temporal context) in which the composition is deployed? Can we make multi-level composition seemingly work as well?

We refer to this problem as the self-composition problem for service ecosystems, where by "self-" we mean that the underlying middleware makes sure that services tend to compose in order to form meaningful new services, in a way that users of the ecosystems simply perceive only atomic and composite services that are actually useful in practice.

We argue that this vision cannot be fulfilled by the above-mentioned approaches taken in isolation, but instead require them to work altogether. Advanced (semantic) matching of services is key to decide whether two services deployed by third parties can be composed together. Evolutionary approaches seem the only solution to the problem of finely tuning the parameters used in the composition, preselecting, for instance, a subset of more promising compositions. And

finally, competition of different compositions based on their actual exploitation appears instead as the only viable approach to measure the quality of a composition, hence promoting successful ones. On the other hand, understanding how they can work together in a comprehensive framework is very difficult, and generally depends on the specific platform and service model one adopts.

16.3.1 Self-Composition with Gradient service

To ground our discussion of requirements and possible solutions, we here outline a paradigmatic example of composition in service ecosystems. We consider a crowd steering scenario, based on the idea of guiding people towards locations hosting events of interest in a complex and dynamic environment (using semantic matching with people's interests) without any supervision, namely, in a self-organised way. In particular, we consider a museum with a set of rooms connected by corridors, whose floor is covered with a network of computational devices (called sensor nodes). These devices exchange information with each other based on proximity, have sensors of various kinds, and hold information as LSAs (e.g., about exhibits currently active in the museum). Visitors exploring the museum are equipped with a smartphone device that holds their preferences. By interaction with sensor nodes, a visitor can be guided towards rooms with a target matching his/her interest, thanks to signs dynamically appearing on the smartphone or on public displays. This system is based on a gradient service, a key brick of self-organisation mechanisms (see Figure 3.1) that is typically designed to provide optimal paths to roam a distributed system even in the context of very articulated environments (a building with rooms and corridors, or a traffic scenario) and by dynamically adapting (namely, self-healing) to unpredicted situations such as sudden road interruption [FMDMSM$^+$12b, FMDMSM12a]. Gradient starts from a source LSA located at an exhibition and spreads clones of it to all the nodes of the network by eco-laws of spreading and aggregation, basic building patterns as shown in Figure 3.1. In turn, as the gradient is stabilised, in each node the minimum distance value is kept, so that the shortest path to the source is automatically (and dynamically) defined [PMV11]. In this way, a person willing to reach a point of interest simply "binds" to a gradient LSA, and public/private displays show signs to guide her to the source.

However, other services exist in the system, such as those provided by the many sensors available around. Among them we have the crowd detection service, enacted by the injection in each node on the floor of a crowd LSA reporting the crowd level sensed there (0 is no crowd, 1 is maximum crowd). An other service is the one that, given visitor preferences and profile, defines a set of other possible points (exhibitions, toilets, refresh points) the user can found interesting or useful. This service can be an external agent that is able to process visitor information, for instance defining related exhibitions, so as to provide, in the form of an LSA, a list of best points the user can pass by. Then we can have sensors for temperature, sensors for pollution, accelerometers etc. More generally, being a pervasive system an open system, new services may appear in the system, and autonomously be composed with existing ones, without a predefined schema. Only those compositions that create useful services will then survive. To exemplify the

idea, in the following we describe two self-composition problem we envisioned.

Crowd composition

This self-composition problem focuses on making the gradient service automatically compose to the crowd detection service so that the estimated distance to the source gets "penalised" in case of a high crowd value, which implies that the path computed towards a point of interest intrinsically takes into account situations of crowded areas dynamically emerging in the environment, properly circumventing them. In particular what we expect to the middleware to automatically do is: decide to compose the gradient with crowd sensing instead of with other (irrelevant sensors), and optimally tune how much crowd information should affect the computation of optimal path.

Ad-hoc path composition

If the gradient service is tagged with the point of interests crossed, it can be composed with the list of "best points" defined for each user, so that, in case of intersection, such gradient gets "reinforced" for instance diminishing its value, which then implies that the direction it comes from has more chance to be chosen, and paths which cross points interesting or useful for the user are likely to be followed. In particular what we expect to the middleware to automatically do is: decide to compose the gradient with the LSA representing the "best points", and optimally tune to which extent information on preferences should affect the computation of optimal path.

16.3.2 A prototype solution

Starting from a number of services available in those nodes where sources of gradient services are located – all reified by an LSA that can be injected in different parts of the system but via chemotaxis reach gradient sources – the solution we propose is based on the following mechanisms:

Composition

Based on one or more "composition recommender" agents available in source nodes, all the available compositions are computed using techniques of advanced matching of interface/behaviour/ontology. These are ranked – according for instance to feedbacks from the users on the composite service, as discussed later – and the description of a limited number of them is reified in the space in the form of a set of LSAs. In the example of crowd steering proposed above, these compositions should result in a set of LSAs representing the source of different gradients, each one created composing the basic gradient with one or more services. We assume that the composition actuated, *i.e.* which services are composed and which parameters are used in the composition, is stored into a property of the gradient LSA, so that the same composition can be performed in all the locations of the system.

Contextualisation

From the source, composite gradients and the basic one, are diffused. In each location they are contextualised with the local value of those services they are composed by. Contextualisations involving gradients normally result in a modification of the local value of distance according to a function δ specified in the service or defined by the composition agent. If services are not locally available, contextualisation does not happen and the distance value does not change. For instance, contextualise gradient with the crowd level perceived in each location at the purpose of discouraging paths crossing crowded areas. It means that distance should be increased and function δ looks like $\delta = k * c$ where k is a function parameter (see later how to define its values), and c is the local crowd level.

Choice

The resulting effect is that the gradient service is available in the system in different forms (the basic one and the composite ones), and users, once entering in the system, can probabilistically choose among them, grounding such decision mainly on the distance and satisfaction values. They will then follow the same composition for the whole path.

Feedback

Each gradient owns a property called "satisfaction", representing the feedback of users on the quality of the service itself. It can for instance evaluate the length of the path, the interest and so on: the higher the value is, the higher is the user satisfaction. People, once they reach the source, should then provide their feedbacks in the form of LSAs with predefined properties —such as length, comfort, beauty, interest. Marks, for each property, might be defined in a predefined range. From there one agent can then be in charge to compute the satisfaction value as a function of all the marks given – for instance, the average. The LSA of the correspondent gradient source is then modified by adding to the previous value of the satisfaction property the last computed. This updated value is then spread in the whole system autonomously as a result of gradient diffusion.

Evaporation

The satisfaction value can also be interpreted as the relevance value of the Evaporation pattern [FMDMSM$^+$12b]. This parameter gets decreased until fading if not augmented by user positive feedbacks. Once it is equal to zero the composite service can be removed. This mechanism ensures that no satisfactory compositions are decayed.

Evolution

If parameters are to be chosen for the composition – for instance parameter k in $\delta = k * c$, representing how much the crowd level has to influence the distance value of the field –, they

can be tuned by agents implementing evolutionary techniques. They are in charge to define the new population of parameters according to the fitness function that can be computed from the satisfaction value of the gradient itself.
The emergent idea is that, after a transition period, system makes available only those services that better fit user preferences. For instance, for those users interested in quickly reach the sources, only those gradients that ensure the lowest arrival time should survive.

16.4 A model for gradients self-compositions

To describe how self-composition of gradients can be built in the SAPERE framework, we first introduce the implementation schema for gradients which we shall rely upon to formalise the self-composition task in the specific example of crowd service exploitation.

16.4.1 The gradient service

Our strategy for modelling gradient's features:

1. a continuous broadcast of information in each node [Bea09], with latter information always overwriting previous one;

2. propagation of estimated distance, computed by summing the contributions given by property `mid:distance` in neighbour annotation;

3. propagation with no horizon (it is easy to impose a distance bound to the propagation of information);

4. after been sent to a neighbour, an annotation is subject to a contextualisation phase (in which e.g. aggregations are executed) before being spread again.

Such rules are reported in Figure 16.1 and are described in turn. Note that they refer to namespace `mid` for URIs related to middleware activities, and to `sos` for those concerning self-organisation aspects.

Rule [PUMP] is a transformation rule which takes a source annotation and creates a new annotation with distance set to 0, used to start the spreading process. The source annotation should feature property `sos:type` assigned to value `sos:source` (we shall say the annotation is of type `source` for brevity), and it should also have some value assigned to properties `sos:aggr_prop` (the name of the property holding the value used to aggregate other annotations as described below), `sos:step` (an incremental value used to refresh information), `sos:r_diff` (diffusion rate) and `sos:r_ctx` (contextualisation rate). This rules fires at diffusion rate `?R`: according to CTMC semantics, this means that as soon as a matching set of reactant annotations is found, actual application of the rule is delayed of a time t randomly drawn according to a negative exponential distribution of probability with average value $1/$`?R`—namely,

Transition Rules for Gradients

```
# [PUMP] An annotation of type source continuously injects the initial
# gradient annotation
?SOURCE sos:type sos:source; sos:aggr_prop ?P; sos:step ?T;
     sos:r_diff ?R; sos:r_ctx ?RC
--?R-->
?SOURCE sos:step =(?T+1) +
?GRAD(?GRAD clones ?SOURCE) sos:type -sos:source sos:diff
     sos:aggr; sos:dist "0"; sos:orientation "here"

# [DIFF] A gradient annotation is cloned in a neighbour, with distance
# increased and updated orientation
?GRAD sos:type sos:diff; sos:dist ?D; sos:r_diff ?R +
?NEIGH mid:type mid:#neigh; mid:remote ?L; mid:orientation ?O;
     mid:distance ?D2
--?R-->
?GRAD + ?NEIGH +
?GRAD1(?GRAD1 clones ?GRAD)  sos:type -sos:diff sos:ctx;
     sos:dist=(?D+?D2); sos:orientation =?O; mid:#loc ?L

# [CTX] A contextualising annotation is transformed back into an annotation
# to be diffused
?GRAD sos:type sos:ctx; sos:r_ctx ?RC
--?RC->
?GRAD sos:type sos:-ctx sos:diff;

# [YOUNGEST] Of two annotations to be aggregated based on property ?P, the
# one with newest information is kept
?ANN1 sos:type sos:aggr; sos:aggr_prop ?P; ?P =[?C]; sos:step ?T +
?ANN2 sos:type sos:aggr; sos:aggr_prop ?P2; ?P2 =[?C]; sos:step ?T2(?T2<?T)
--->
?ANN1

# [SHORTEST] Of two annotations to be aggregated based on property ?P,
# the one with shortest distance from source is  kept
?ANN1 sos:type sos:aggr; sos:aggr_prop ?P; ?P =[?C]; sos:dist ?D1;
     sos:step ?T  +
?ANN2 sos:type sos:aggr; sos:aggr_prop ?P2; ?P2 =[?C];
     sos:dist ?D2(?D2>=?D1); sos:step ?T
--->
?ANN1

[DECAY] An annotation decays
?GRA sos:type sos:diff; sos:r_dec ?RD
--?RD->
```

Figure 16.1: Rules for gradients

it is fired at frequency `?R`. The effect of firing this rule is twofold: it increases the value of property `sos:step` of the source annotation, and creates (by cloning the source) a new annotation `?GRAD` of type `sos:diff` and `sos:aggr` (original type `sos:source` is removed), and declaring distance 0 and orientation `here` with respect to the source—these properties will be update as the annotation is spread around. Note that the right-hand side of a rule mentions only changes in annotations, without any need of repeating information provided in the left-hand side which has not changed: this allows a simpler structuring of rules with respect to more syntactic approaches [PMV11].

Rule [DIFF] is used to diffuse a cloned version `?GRAD1` of a gradient annotation `?GRAD` in one neighbour. The gradient annotation should be of type `diff`, declare distance `?D` from the source and have diffusion rate `?R`. This rule has as further reactant a neighbour annotation, from which information about a neighbour `?L` with orientation `?O` and distance `?D2` can be extracted. This rule leaves the reactants unchanged, but creates a clone of `?GRAD` with type `ctx` instead of `diff`, to be relocated at `?L` and indicating orientation (with respect to originating location) `?O` and distance `?D+?D2`. Note that continuous application of this rule at rate `?R` makes copies of the gradient annotation to be created and diffused in all neighbours.

An annotation is relocated with type `ctx`, which forbids further application of rule [DIFF]. This is because we first need to aggregate all incoming annotations before a new diffusion is scheduled. To do so, rule [CTX] defers change of type from `ctx` to `diff`, which happens at rate `?RC` (contextualisation rate).

One form of aggregation is due to rule [YOUNGEST], used to refresh gradients with new information. It takes two annotations such that the values associated to the aggregation property `?P` are the same, and retains the one with bigger `sos:step`. The idea is that `sos:aggr_prop` holds the name of a property `?P` which is expected to contain some value(s) that should be identical in two annotations for them to be aggregated. For instance, it could be the unique id of a gradient source, so that we do not aggregate annotations coming from difference sources—but more involved situations can be programmed as exemplified in next section.

Rule [SHORTEST] is similar: it takes two annotations with same step (hence coming from the same gradient annotation as generated by rule [PUMP]), and retains the one with shortest distance only if, again, they have same content of aggregation property.

Finally, rule [DECAY] is used to remove an annotation at a given decay rate `?RD`. This is useful as a cleanup mechanism in the case a gradient source is removed from the system—though it plays no crucial role in and will be neglected in next discussions.

For the above rules to properly work we need to synthesise a well-structured source annotation, featuring all required properties and proper values for rates—namely, diffusion rate is to be chosen to keep the system sufficiently reactive to changes without bloating the system with undesired messages, and contextualisation rate should be significantly higher than diffusion rate (at least one order of magnitude). For instance we could use the following source annotation:

```
:id314 mid:#loc :loc117;  sos:type sos:source;
       sos:step "0";      sos:sourceid "341AB2"
```

Eco-laws for gradient contextualisation

```
# [COMPOSITION] The gradient source is composed with the crowd service
?SOURCE sos:type sos:source; scm:satisfaction ?S +
?CROWD scm:type crowd; crowd:level ?CL
--->
?SOURCE +
?CSOURCE(?CSOURCE clones ?SOURCE) scm:property sos:dist;
     scm:parameters scm:crowd_op  ?CF;  scm:crowd_op  ?CF*?CL

# [CONTEXTUALISATION] If sensors perceive crowd, the gradient distance is
# augmented
?GRAD sos:type sos ctx; sos:dist ?D;  scm:property sos:dist;
     scm:parameters scm:crowd_op scm:crowd_factor; scm:crowd_factor ?CF;
     scm:crowd_op  ?CF*?CL +
?CROWD scm:type crowd;  crowd:level ?CL
--->
?CROWD + ?GRAD  sos:type -sos:ctx sos:diff; sos:dist =(?D+?CF*?CL)

# [FEEDBACK] Feedbacks are used to update the satisfaction values
?FEEDBACK scm:parameters scm:crowd_op; scm:feedback  scm:velocity;
scm:velocity ?V +
?GRAD scm:satisfaction ?S; scm:parameters scm:crowd_op
--->
?GRAD scm:satisfaction =(?S+?V)

# [EVAPORATION] The gradient satisfaction value gets decreased
?GRAD scm:satisfaction ?S; scm:factor_ev ?FE; scm:r_ev ?RE
--?RE->
?GRAD scm:satisfaction =(?FE*?S)

# [DECAY] If the gradient satisfaction value becomes zero that composition
# is removed
?GRAD scm:satisfaction "0";
--->
```

Figure 16.2: Additional rules for anticipative gradients.

```
sos:aggr_prop sos:sourceid;
sos:r_diff "10";   sos:r_ctx "100"
```

16.4.2 Rules for Gradient Composition

Rule [COMPOSITION] in Figure 16.2, is in charge of creating the source of the composite gradient. The new source owns the same set of properties of the basic one, as soon as it refers to the same PoI, but it also defines which properties are to be modified as an effect of the composition – `sos:dist` –, and which are composition parameters – `scm:parameters` – ,

i.e. function – `scm:crowd_op` – to be used for modifying the distance value, and coefficient —`scm:crowd_factor`. In this preliminary example we adopt a simple linear function of the crowd level perceived by sensors in each location. It increases the distance of those paths crossing crowded area, if the coefficient has an opportune value. Rule [CONTEXTUALISATION] applies the composition specified in `scm:parameters`, in each location of the system. Rule [FEEDBACK] is used to modify the satisfaction value of gradient, taking into account the velocity of the path chosen. It can be computed by the system itself, or provided by the user. The value of `scm:satisfaction` gets increased by `scm:velocity` value.

Rules [EVAPORATION] and [DECAY] are finally used to evaporate the `scm:satisfaction` value according to an evaporation factor `scm:factor_ev ?FE` in the range of [0,1] and to remove the gradient composition in case of a `scm:satisfaction` value equal to zero, or negative, ensuring that only those compositions that compute fast paths will survive.

16.5 Towards simulation of gradient self-compositions

We checked the correctness of the proposed self-composition solution by simulation, conducted using ALCHEMIST.

Results presented show early experiments on gradient composition with crowd level, where people dynamically enter in the system and begin to move towards the target ascending one of the different compositions available. Movement velocity is constant in the system, except for crowded areas, where it decreases according to the crowd level perceived: higher crowd, slower velocity. In these early experiments, parameter k in function δ is not computed dynamically through an evolutionary algorithm but composition agents define a set of possible values. A new composition corresponds at each value. Users chose one gradients randomly but considering also the associated satisfaction value that measures the velocity of the best path each gradient proposes. This feedback is computed by each user once reached the target, given the length of the path and the distance walked.

The goal is to observe that satisfaction values, initially the same for each composition, dynamically change according to feedbacks from the users, and only the best compositions survive. In Figure 16.3 such process is shown for different values of parameter k: the system reaches a stable state with $k = 0.5$ ensuring the faster path while the other compositions decay.

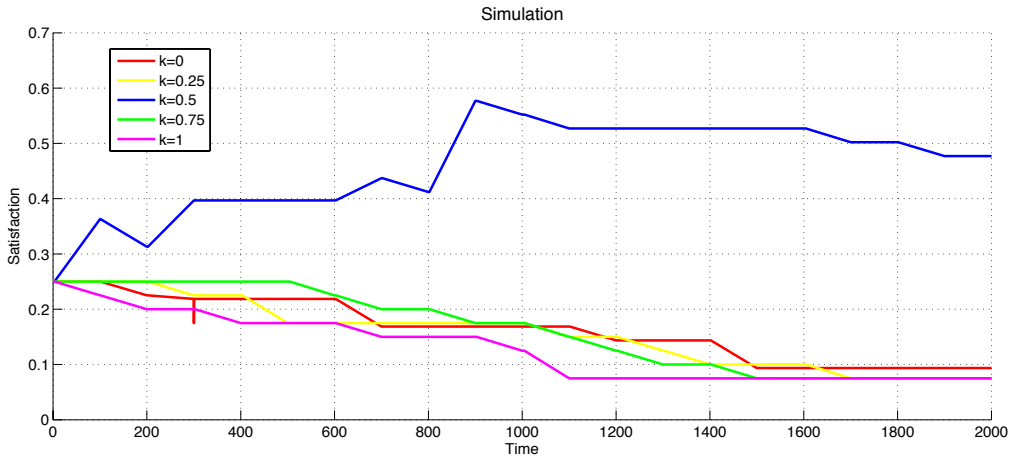

Figure 16.3: Satisfaction values for different compositions changing over time. k is the parameter that differentiates services. The system stabilises with $k = 0.5$.

17

Case study: Semantic resource discovery

Supporting expressive resource discovery atop this unstructured, highly-dynamic substrate presents a challenge, with the combination of semantic reasoning, context-awareness and self-organisation techniques offering intriguing possibilities: semantic matching provides a flexible framework for quantitatively assessing the suitability of resources to a given request; context-awareness may account for application quality of service concerns that are important, but orthogonal to the matching process, and self-organisation techniques provide a means of adapting to changing environmental conditions—the arrival and removal or resources, the mobility of interacting agents, and the ever evolving network topology.

To illustrate the utility of this approach, we consider a smartphone application for an exhibition centre (for example, Tokyo's Makuhari Messe, which contains eleven exhibition halls and has a footprint of $210,000m^2$), which offers routing information and exhibit suggestions based on a user's profiled interests. These techniques support the development of exploration strategies to:

- select the exhibit most suited to the user's interests;
- select preferential exhibits while accounting for factors such as their distance from the user's position or crowdedness along the route;
- route the user to an exhibition hall containing many exhibits of interest, but not necessarily the most interesting if it is considered spatially isolated;
- use a combination of the above strategies, for example, to guide the user to a particular exhibition hall and then route them to the exhibits inside the hall in order of interest or least crowded.

Match Degree	Description	DL Notation	Similarity (R, A)
Exact	The request and advertisement are equivalent concepts.	$R \equiv A$	1
Subsume	The request expresses a more general concept than the advertisement.	$A \sqsubseteq R$	1
Plugin	The request expresses a more specific concept than the advertisement.	$R \sqsubseteq A$	$\frac{\lvert S(A)\rvert}{\lvert S(R)\rvert}$
Intersection	The intersection of the request and advertisement is satisfiable.	$\neg(R \sqcap A \sqsubseteq \bot)$	$\frac{\lvert S(A) \sqcap S(R)\rvert}{\lvert S(R)\rvert}$
Disjoint	The request cannot be satisfied by the advertisement.	$R \sqcap A \sqsubseteq \bot$	0

Table 17.1: Match-degrees and their semantic scores, defined using the similarity metric proposed by Skoutas et al. [SSS07]. $S(X)$ denotes the set of super-concepts of concept X in its defining ontology.

17.1 Semantic Matching in SAPERE

17.1.1 Semantic Matching

A matchmaking process takes a request and a set of resource descriptions as input, and outputs the set of resources that satisfy the request. Many matchmakers adopt simple syntactic schemes; for example, named interfaces or predefined categories [Wal00], or sets of attribute-value pairs to which string and numeric comparisons can be applied [GPVD99]. Such approaches are limited by an inability to compare semantically equivalent but syntactically different concepts or handle approximation. To illustrate, consider the concept hierarchy depicted in Figure 17.1, based on the Vehicle Sales Ontology[1]. Here, the concepts `Automobile` and `Car` may be considered as equivalent terms, or `TwoDoorCar` as an approximate match to a `FourDoorCar`. Semantic analysis supports the incorporation of such matches where they would otherwise be discounted.

A semantic matching process has three parts: a model for describing requests and resources, a scheme for measuring the semantic similarity of concepts, and an algorithm for evaluating the match.

Modelling Resources and Requests

A request expresses the intersection of a number of properties, each named concept or existential restriction belonging to a resource. Using Description Logics (DL) notation [BCM+03], the following describes a request for two door, electric cars manufactured after 2009, and three advertisements for candidate vehicles to be matched against this request.

[1] `http://www.heppnetz.de/ontologies/vso/ns`

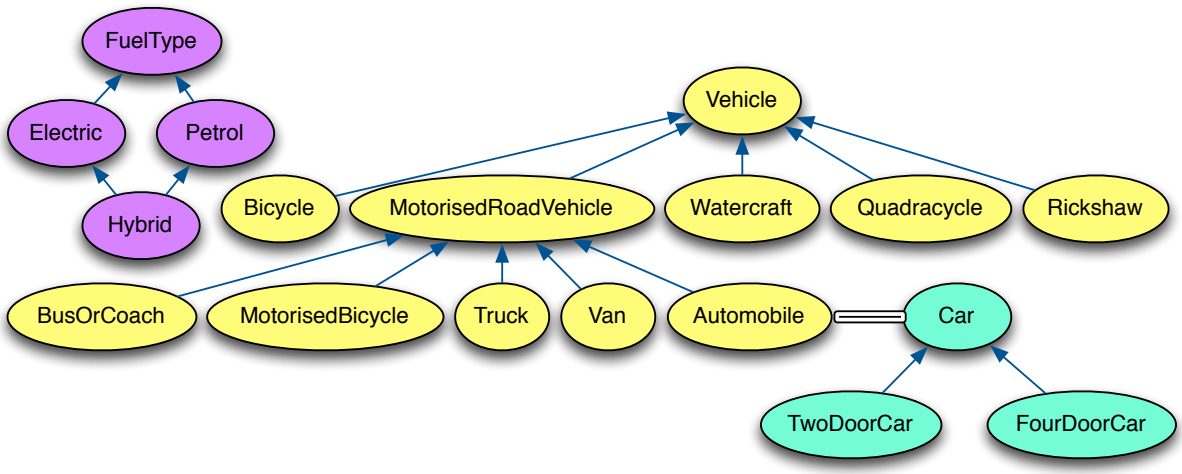

Figure 17.1: A partial illustration of concepts in the Vehicle Sales Ontology, an equivalence relation from one term to another concept hierarchy, and a snippet from a fuel type ontology.

R: $TwoDoorCar \sqcap \exists modelDate \geq 2009 \sqcap \exists fuelType = Electric$
$A1$: $Automobile \sqcap modelDate = 2009 \sqcap fuelType = Hybrid$
$A2$: $TwoDoorCar \sqcap modelDate = 2011 \sqcap fuelType = Petrol$
$A3$: $Van \sqcap modelDate = 2011 \sqcap fuelType = Electric$
$A4$: $Bicycle \sqcap modelDate = 2007$

None of the four advertisements fully satisfies the request: there is insufficient information to determine if $A1$ is a $TwoDoorCar$; $A2$ differs in the requested fuel type; although electric, $A3$ is a Van; and $A4$ satisfies no request aspects. However, by inspection we can see that some of the advertisements more closely relate to the request than the others. We say that an advertisement is *compatible* with a request if the intersection of their descriptions is satisfiable.

Quantifying Semantic Similarity

[PKPS02] and [LH03] introduce ordered degrees of matching to categorise the semantic compatibility between a request (R) and an advertisement (A), of which there are five: *exact*, *plugin*, *subsume*, *intersection*, and *disjoint*. [BPd$^+$08] extends this model with scoring mechanisms to support differentiation of matches falling within the same category. Descriptions of the match-degrees and formulas for scoring taxonomically related concepts are shown in Table 17.1.

Applying this mechanism to the concept hierarchy in Figure 17.1 allows us to quantify the similarity between terms in our earlier example. For example, the degree of similarity between $Automobile$ and $TwoDoorCar$ is 0.75 (plugin), the degree of similarity between $TwoDoorCar$ and Van is 0.5 (intersection), and the degree of similarity between $Hybrid$ and $Electric$ is 1 (subsumption).

Executing the Semantic Match

A semantic matching algorithm conflates all the terms in a request and resource description and may be realised in many ways.

A possible implementation of a `MATCH` algorithm takes two variables R and A as input, which correspond to the request and advertisement or some sub-set of their terms respectively. The function first checks if R subsumes A; if so, the algorithm terminates with a similarity score of 1. If R and A refer to atomic named concepts or literals, the similarity score is calculated by a call to the `SIMILARITY` function, which resolves the comparison to a score depending on how the concepts are related. If not, R and A are composites consisting of named concepts or existential restrictions. This case is handled by recursively decomposing R into its constituent parts, $R_{1..n}$, and averaging the maximum score for each R_i, when compared with corresponding parts of A, that is, $A_{1..m}$. Applying this algorithm to the above example yields the scores: $A1$=.92, $A2$=.83, $A3$=.83, and $A4$=.08.

17.2 Distributed Resource Discovery

We now clarify how standard self-organisation patterns can be augmented with semantic reasoning so as to support a decentralised approach to resource discovery in pervasive computing applications.

17.2.1 Self-organisation Patterns

This section briefly reviews those patterns that form our interaction building blocks.

Spreading

The Spreading pattern progressively sends information over the system from one node to its neighbourhood, iteratively, so as to make it available globally by using only local interactions. In our framework it is supported by the following eco-law:

Spreading

```
?A :val ?V; :diff ?F; :loc ?L
-->
?A + ?B :#clones ?A; :val =?W; :loc ="*"; :prev = ?L
?BIND (:exec2(?F,?V) AS ?W)
```

Aggregation

The Aggregation pattern reduces the amount of information in the system, typically by summarising data disseminated by agents over time, e.g., by spreading. It is supported by:

Aggregation

```
?A :src ?S; :aggr ?F; :val ?V + ?B :src ?S; :val ?W
-->
?A :val =?Z :
?BIND (:exec3(?F,?V,?W) AS ?Z)
```

Decay

The Decay pattern incrementally ages information in the system. A monotonically decreasing function is applied to a numeric decay value over time, with the information removed upon the counter reaching zero. Decay-based Aggregation supports the overwriting of old information from a data source with the more recent. Decay is supported by the following eco-laws:

Decay

```
?A :decay ?D :dfunc ?F;
-->
?A :decay ?N;
?BIND (:exec4(?F,?D) AS ?N)

?A :decay 0;
-->
```

Gradient

The Gradient pattern is a composition of the Spreading and Aggregation patterns, where information about the sender distance and direction is propagated across an extent of the network. An LSA can become source of a gradient by specifying initial distance 0, a diffusion function incrementing the distance value, and an aggregation function retaining the less traveled LSA among competing gradient LSAs from the same source [VPMS12]. The Decay pattern supports *dynamic* gradients by discarding old information as a gradient source moves within the network.

Chemotaxis

The Chemotaxis pattern provides a mechanism to route information towards the source of a gradient via the shortest accessible path. In our framework, movement of LSAs towards a gradient source is achieved by diffusing in the direction indicated by property `:prev` as created by the Spreading eco-law.

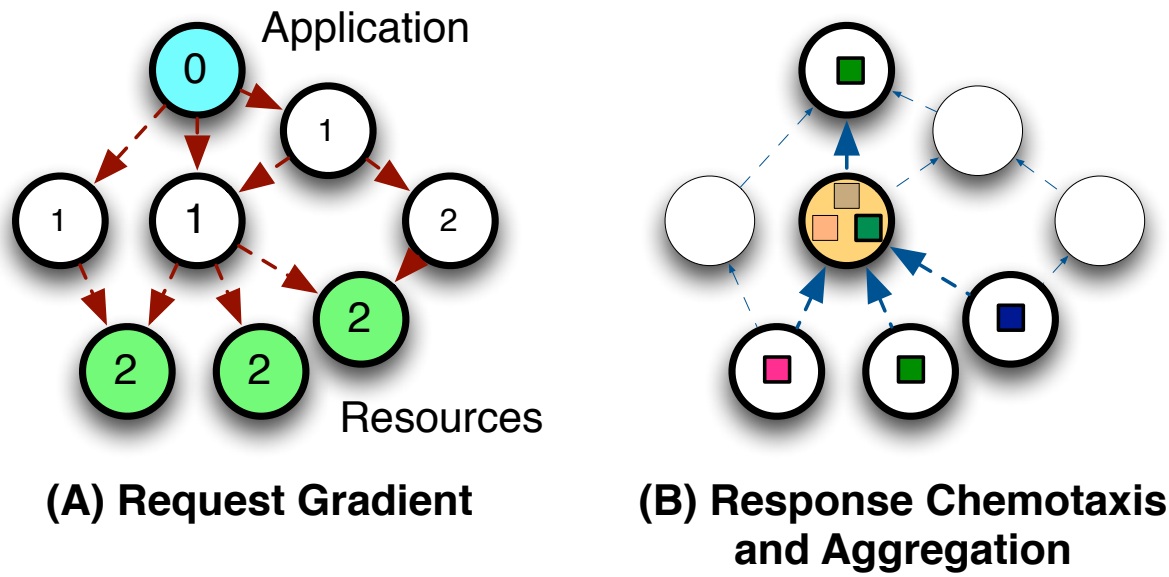

Figure 17.2: Resource discovery interactions: (A) a request gradient is spread, and (B) replies ascend the gradient and aggregate.

17.2.2 Self-organising, Context- and Semantics-aware Resource Discovery

Building upon the above concepts, we now introduce the key steps in our approach to self-organising, context- and semantics-aware resource discovery, which are illustrated in Figure 17.2.

Establish the resource request

An application searching for a resource publishes a `request` gradient LSA carrying information about the source *R* of the request: spreading and aggregation eco-laws will iteratively fire in random establishing a gradient data structure, with horizon *H* (hop distance counter), and in which each LSA has a pointer to the node in which it was created, which is also the node indicating direction to the source following shortest path.

Resolve matches and publish reply

Where a gradient LSA containing a request is co-located with an annotation describing a resource (or service) *S*, we apply semantic matching, as described in Section 17.1 to resolve their compatibility to a match-degree. This is realised by a matching eco-law that creates a `reply` LSA pointing towards the request and the service, and with the value *MD* representing the match-degree (greater than 0).

Route match information to the requestor

The reply LSA is created such that it diffuses by copying towards the request source; namely, the chemotaxis pattern is applied to deliver a remote representation of the resource annotation matched to the request source. This is continually applied to maintain a live link between the two points in the network: the request site and service site. This ends after the request gradient expires, for example, after the request is explicitly removed. This allows the process to update and self-heal under conditions where:

- an intermediate node used for routing fails;
- a compatible resource appears within the gradient;
- the requestor or resources are mobile within the network.

Augmenting matches with application context

As the reply LSA is routed, aggregation may be employed to modify, either positively or negatively, the reported match-degree, using information external to the reply. The approach of aggregating LSAs contextually to routing was inspired by the work of Fernandez-Marquez et al. [FMST$^+$12]. This is achieved by applying an aggregation function to any LSA directed towards a request source that accounts for contextual factors (*any information that can be used to characterise the situation of entities that are considered relevant to the interaction* [DAS01]) specified by the application; and with such information readily available for consumption through its expression in an LSA. For example, $\langle S_0, 0.9, 0.75 \rangle$ (resource, original-match-degree, match-degree) might represent a reply LSA that has been updated to account for the distance an LSA has travelled, or the presence of physically crowded regions along the response path.

Aggregate results en route by competition

Where a request is matched to multiple, distributed resources, the algorithm returns the most suitable to the requestor. For example, in the case where there are two co-located response annotations of the form $\langle S_0, 0.9, 0.75 \rangle$ and $\langle S_1, 0.8, 0.8 \rangle$ (resource, original-match-degree, match-degree), the latter will be propagated while the former will be dropped. Combining this with the previous mechanism provides a contextual-semantic approach whereby a resource may be retained that, although it has worse match-degree when considered in isolation, better accounts for a requestor's contextual requirements.

Aggregate results en route by collaboration

An alternative to the above approach is to augment the standard semantic-matching process by making a resource more, or less, attractive depending on some contextual relation to similarly matched resources. For instance, in applications where resources represent physical points of

interest (POIs), we may consider a small region of the space that includes many POIs more favourable than another region containing the POI with highest match-degree. This is achieved by an aggregation function that "joins" two replies instead of selecting one of them. For example, response annotations of the form $\langle S_0, 1, 0.9 \rangle$ and $\langle S_1, 2, 0.75 \rangle$ (resource, distance, match-degree), can be subject to joining as they represent spatially close services. Hence, they are aggregated into a single annotation $\langle S_0 + S_1, 1.5, 0.85 \rangle$, in which distance and match-degree summarise the original ones. As before, application specified contextual factors may readily be accounted for.

Receipt of replies

The requestor will eventually receive aggregated replies, providing information about (summary) distance, (summary) match-degree, resource provider (or resources provider), and direction to reach the service. Such information can then be used, for example, to steer the requestor agent towards a selected resource or selected resources [VPMS12].

Dynamic self-healing

The evolution of ecosystem states over time is supported by the Decay pattern, which is employed to periodically remove older information while newer information is concurrently refreshed. The periodic update of request gradients and reassertion of responses accounts for the requestor moving closer to or further away from resources, for resources disappearing or arriving, and for the corresponding impact these changes have on the aggregation applied to the match degree results in transit.

17.3 Eco-laws to support matching

In this section we present eco-laws that support the contextual-semantic management of resource discovery, describing their behaviour and the functions they encapsulate. For the sake of space we will not provide the general-purpose eco-laws supporting the spreading, aggregation, gradient, and chemotaxis self-organisation patterns, for they are discussed elsewhere [VPMS12]. We instead focus on the four eco-laws that incorporate the key strategies described in previous section, namely, the generation of replies along with a match-degree, their contextualisation at each node they transit across, and the competition-/collaboration-based aggregation as they are routed towards the requestor. They are presented in Figure 17.3, and described in turn. These eco-laws make use of a default (omitted) name-space corresponding to the ontology of resource discovery, the `sos` namespace incorporating the concepts relating to self-organisation patterns, and the `sapere` namespace incorporating all the general concepts of pervasive service ecosystems.

Eco-law `[MATCH]` creates the reply. It takes a request LSA `REQ` and a service LSA `SER` located in the same node, and creates the reply LSA `REP` such that: it clones `REQ`, has type `:reply`, points (via so-called bonds) to the service LSAs (`?SER`) and stores its content (`?S`),

Resource discovery with match-degree and spatial information

```
% [MATCH] When a request and service match, a reply LSA is generated to be diffused towards
% requestor
?REQ :type :request; :content ?R; :source ?SRC + ?SER :type :service; :content ?S
--?Rate-->
?REQ + ?SER +
?REP :#clones ?REQ; :type = :reply; :service ?SER; :content ?S; :original_match_degree ?MD;
     :match_degree ?MD; :ctxed "false"; sos:direction ?SRC; sos:switcher :ctxed;
     sos:type sos:ascend
BIND(:semanticMatchDegree(?R, ?S) AS ?MD)
FILTER(?MD > 0)
BIND(:replyRate() AS ?Rate)

%[CTX] A reply contextualises, and is accordingly propagated one-step
?REP :type : reply; :ctxed "false"; :ctxprop ?P; :ctxtype ?T; :ctxfun ?F; ?P ?V
?CTX :type ?T; ?P ?V2
--?Rate-->
?REP :ctxed = "true"; ?P = ?W +
?CTX
BIND (:exec(?F,?V,?V2) AS ?W)
BIND(:diffRate() AS ?Rate)

% [AGGREGATE-COOP] Joins information from two replies belonging to the same service cluser
?A1 :type :reply; :ctxed "true"; :request ?R; :service ?S1; :ctxprop ?P; ?P ?V1;
     :match_degree ?MD1; :coop_pred ?CP; :coop_md ?CM; coop_ctx ?CC +
?A2 :type :reply; :ctxed "true"; :request ?R; :service ?S2; :ctxprop ?P; ?P ?V2;
     :match_degree ?MD2
--->
?A1 :service ?S; ?P = ?V; :match_degree = ?MD
FILTER(:exec(?CP,?V1,?V2)) .                    % The two services should "cooperate"
BIND (:exec(?CC,?V1,?V2) AS ?V) .               % Let V be the aggregated contextualisation
BIND (:exec(?CM,?MD1,?MD2) AS ?MD) .            % Let MD be the aggregated match-degree
BIND (:union(?S1,?S2) AS ?S) .                  % Let S the aggregated service description

% [AGGREGATE-CHOOSE] Selects between two replies originating within different service clusters
?A1 :type :reply; :ctxed "true"; :request ?R; :service ?S1; :ctxprop ?P; ?P ?V1;
     :match_degree ?MD1; :comp_pred ?CP; :coop_sel ?CS +
?A2 :type :reply; :ctxed "true"; :request ?R; :service ?S2; :ctxprop ?P; ?P ?V2;
     :match_degree ?MD2
-->
?A :service ?S1; ?P ?V; :match_degree ?MD1
FILTER(:exec(?CP,?V1,?V2)) .                    % The two services should "compete"
FILTER(:exec(?CS,?V1,?MD1,?V2,?MD2)) .          % Service A1 has a stronger match-context pair
```

Figure 17.3: Eco-laws for handling request-reply match, and for distributed aggregation.

it is also of type `sos:ascend` relative to the gradient generated by `SRC`, and finally specifies properties `:ctxed` as flag to state whether contextualisation already took place. It also describes the match-degree `?MD`, computed by function `:semantic MatchDegree` over service content `?S` and request content `?R`, implemented by the algorithm shown in Section II—only match degrees greater than 0 are realised as a reply. Note that this eco-law is reapplied over time so as to continuously support the chemotaxis pattern as already mentioned, as such, external function (constant) `:replyRate` is used to extract the application rate `?R`.

Eco-law `[CTX]` ensures a reply LSA `?REP` is properly contextualised before being diffused one step. It seeks another LSA of type `?T` having property `?P` assigned to a value `?V2`, and accordingly changes `?REP` by switching the value of `:ctxed` flag and updating `?P` to the result of applying function `?F` to `?V2` and previous value `?V`. Switching such a flag causes the diffusion eco-law to move this LSA to a neighbour in the direction that ascends the request gradient, and to switch the flag back to false—so that eco-law `[CTX]` can be applied again. Before diffusion takes place, however, aggregations should occur according to the following eco-laws.

Eco-law `[AGGREGATE-COOP]` joins two reply LSAs perceived as belonging to a unique cluster of services (e.g., related points of interest). It takes two reply LSAs for the same requestor `?R`, whose contextualisation values make them be perceived as "near" by requestor-specified predicate `:cooperate`, and creates a new reply LSA (overwriting one of the two originals) in which the service indication is obtained by joining `?S1` and `?S2` by external function `:join`, contextualisation value `?V` is computed by function `?CC`, and match-degree is computed by function `?CM`.

The implementation of such functions strictly depends on the kind of contextualisation an application requires. In our simulations in the next section we shall use an estimated physical distance as a contextual property, and as such we consider the following function (to be compared with other approaches in the next activities of this research): `:coop_pred` simply checking that "distances" from the two resource are below a given threshold; `:coop_ctx` the weighted mean of the two distances based on match-degree, namely, $(?V1 * ?MD1 + ?V2 * ?MD2)/(?MD1 + ?MD2)$; and `coop_md` implements a t-conorm (a function to perform a "fuzzy union" [BS11]), and in particular the so-called product logic t-conorm $?MD1 + ?MD2 - ?MD1 * ?MD2$. Finally, `:join` simply appends the two arguments yielding the list of all service identifiers.

Eco-law `[AGGREGATE-CHOOSE]` selects one from two reply LSAs whenever they are perceived as belonging to different clusters of services. Similarly to the previous law, it takes two reply LSAs for the same requestor `?R`, whose contextualisation values are such that they are not perceived as "near" by external function `?CP`, and selects the one with the stronger semantic-spatial situation (discarding the other). This is computed by function `?CS`. In our reference case study: `:comp_pred` is exactly the opposite of `:coop_pred`, while `:comp_sel` checks whether $(?MD1 * k/(?V1 + k)) \geq (?MD2 * k/(?V2 + k))$ for some fixed parameter k. This compares the two match-degrees, but penalises the one having greater distance (e.g., when $V = k$ the actual match-degree is halved).

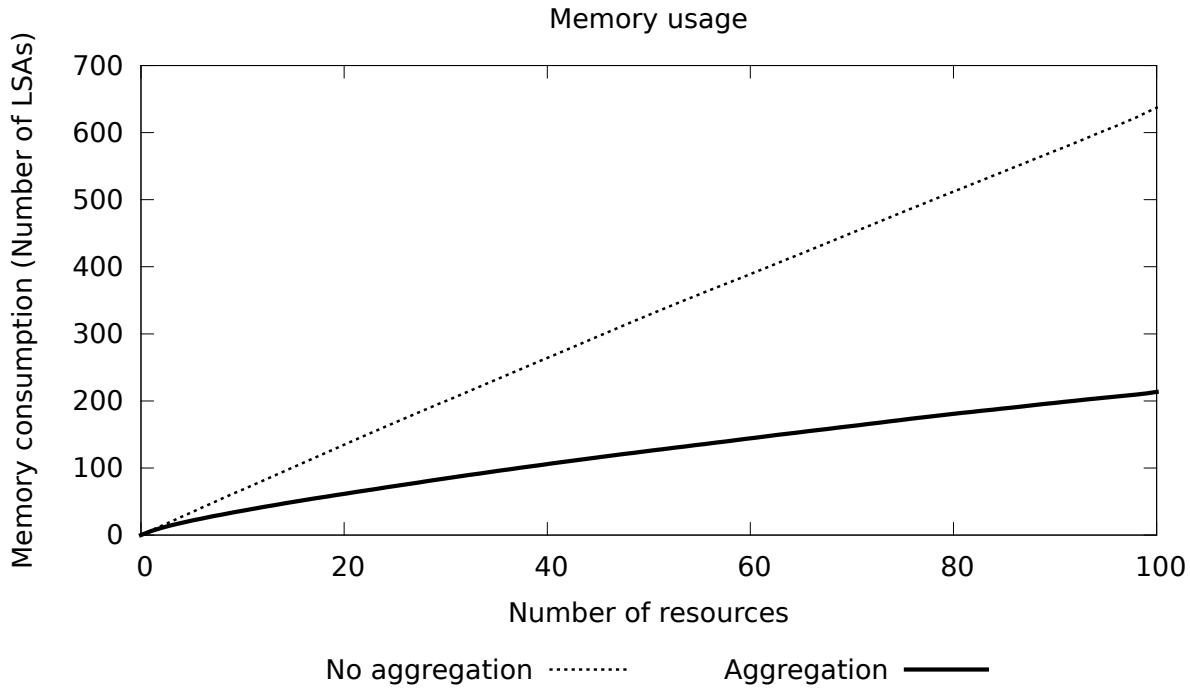

Figure 17.4: Number of reply LSAs depending on the number of matching resources.

17.4 Evaluation of Approach in ALCHEMIST

The simulated scenario aims to provide early validation of the model and functions (matching and aggregation operators) proposed in Section 17.3 to select the best service among those matching a request, and investigate the performance gain due to the use of our distributed aggregation techniques.

As a first step, we measured the overhead advantages due to the use of a distributed form of aggregation, both in space and in bandwidth, which we estimate by considering the number of reply LSAs that stabilise in the distributed set of node as a request is issued. We simulated a bi-dimensional environment populated with an infrastructural grid of 11x11 nodes in a rectangular structure. A request is initially injected from a boundary position. In each simulation we randomly placed N (ranging from 1 to 100) resources in the environment, each with a match degree randomly chosen between 0 and 1. The request generates a gradient, and all resources eventually send a trail of LSAs back towards the requestor. We compared the case in which no aggregation of such replies is performed, with the case in which as soon as two reply LSAs are co-located they compete such that only one of them is retained: it can hence be easily expected that a smaller number of LSAs will ultimately be diffused around. Note that the results of this simulation are independent of the aggregation behaviour. Results are reported in Figure 17.4, each showing an average value out of 50 runs. They show that while the overhead in a standard

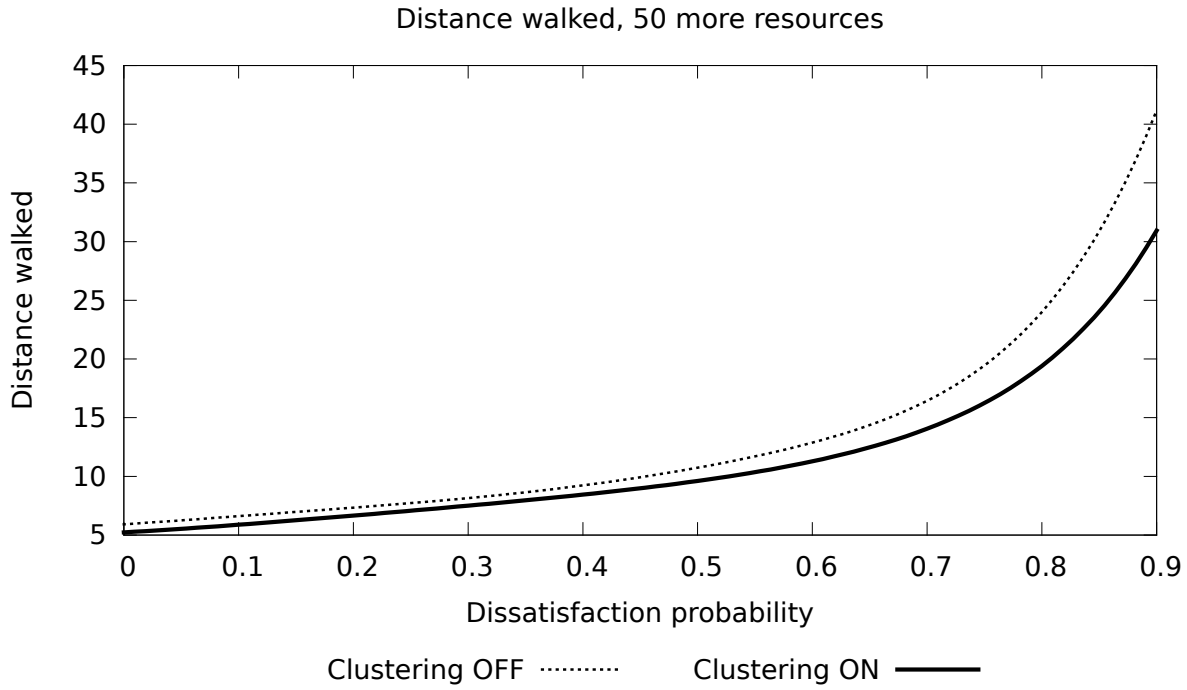

Figure 17.5: Average distance travelled over fault probability

approach clearly grows linearly with the number of matching resources, distributed aggregation makes the number of involved LSAs grow much slowly, generally $2/3$ less for the particular case we considered—of course, different topologies can lead to different levels of improvement.

As a second example we show the advantage of using a cooperative approach during aggregation, which makes clusters of resources be preferable with respect to isolated resources. The reason why this is considered useful is that frequently (for example, think of a general shopping scenario), the first matching resource reached is not the right one to pick: a new resource might be discovered from there—depending on the application scenario, this may happen if the resource is no longer available, or the application user is not satisfied despite the match and wants to quickly find a nearby alternative. In that case a new request has to be issued: if the user already reached a cluster of matching resources, then it will more likely find another match nearby. To simulate the advantages of this behaviour we define a *fault probability* p that, as one resource is selected from the requestor in a node n, another one needs to be searched because of dissatisfaction. Our reference environment is the same as previous case, but we added 100 resources on one side of the rectangular area, to simulate a non-homogeneous displacement of resources that can take advantage of our management of clustering.

Figure 17.5 and Figure 17.6 report results, each considering a different value of p: the former measures average distance travelled (considering that many resources are generally to be reached), the latter measures the overall result quality, which is computed by the same formula

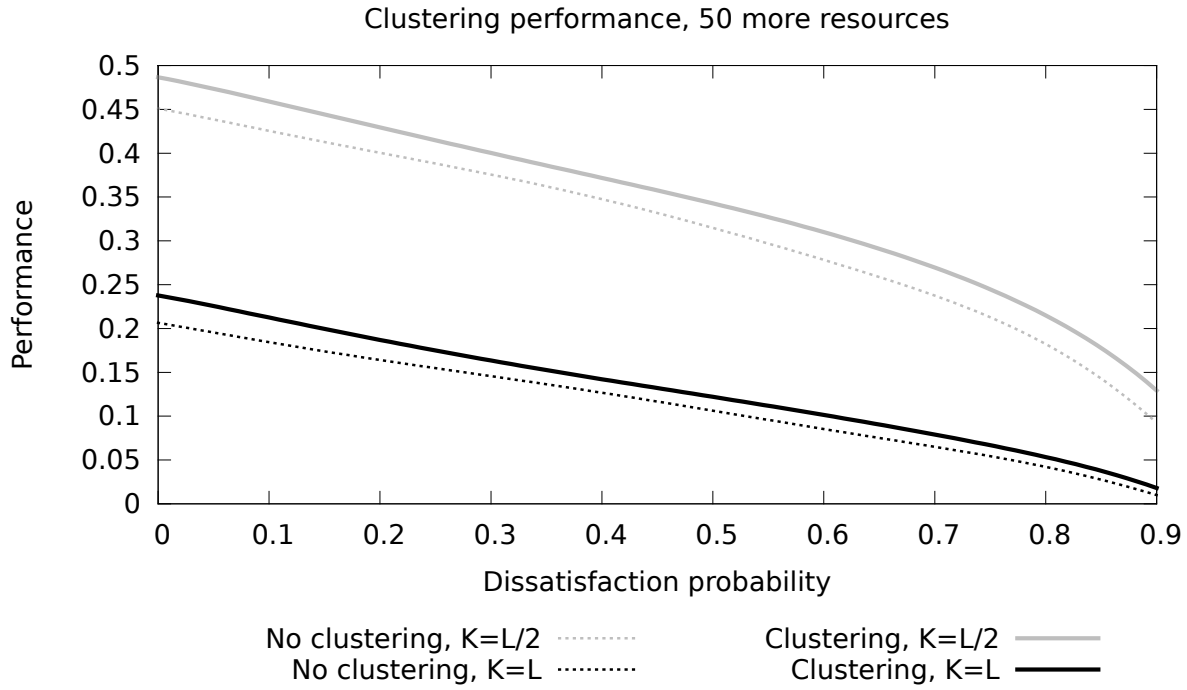

Figure 17.6: (Average match-distance value over fault probability.

$k * M/(D+k)$ we used to compare competing replies, that is, promoting resources which both are good matches and are near to the requestor. Both charts show that, even if the environment has no specific clusters, there is a significant gain in the quality of replies.

18 Case study: Crowd disasters prediction

18.1 Human Crowd Density and Safety

This year, more than 60 people died during the "New Year's Eve stampede" in Abidjan[1]. During the Love Parade on July 24, 2010 in Duisburg, 21 people died and more than 500 were injured in a crowd disaster [HM12]. Similarly at Hillsborough stadium on April 15, 1989, 96 people died and 400 were injured [NR95]. Many more such sobering statistics exist, explaining why disciplines such as crowd density estimation and human movement flow analysis have received much attention over the years. While human movements generally depend on the environment and the current context – such as the type of event – it was shown that disasters occurred in many different venues at different circumstances and could have been avoided through better management and planning [Fru93]. Towards this end, crowd density metrics that indicate the danger for individuals in a crowd have been established [HM12, Fru93]:

- density of people in a defined area, with safe levels approximated as 2-3 persons per meter;
- flow of persons per meter, safe below 82 persons per meter and minute;
- pressure exerted on each person, tolerable up to 1500 N.

Thus, measuring crowd densities and predicting their development can offer valuable information about the current and future overall event safety and help to avoid catastrophes.

Previous work on this topic has been conducted in the field of pedestrian crowd dynamics [HBJW05, BKSZ01, ABW06] and similarly in the field of evacuation scenarios [HJ09, KRS10, ZRFS11]. The purpose of many of these works has been to further our understanding of these dangerous phenomena in order to pre-emptively react to them. Generally the approach of choice has been to use simulations and draw conclusions from their outcome, whereby simulations have been based on models of reality that were constructed by abstracting data from real-world events. Our contribution uses a similar approach, in that we use time series analysis and other approaches

[1] `http://www.bbc.co.uk/news/world-africa-20885814`

to predict the crowd density at future points based on observed crowd densities from the past. Our approach is based on a dataset – comprised of the spatial properties of event spectators – that was created by a collaborative crowd-sourcing/participatory sensing approach using mobile devices of spectators (similar to e. g., [WFR+12]; see [BEH+06] for more information on participatory sensing). It has been argued that a linear correlation between user density – meaning the observed density of spectator's mobile devices – and crowd density exists [WFR+13]. Due to this linear relation a prediction of user density can be assumed to be correlated to actual crowd density. Thus for the purpose of this work we do not distinguish between them. However all presented results related to user density and should be viewed in this light. We state our research hypotheses as follows:

Research Hypotheses:

(H1): Data on crowd movements that was observed during actual mass events can be used to develop prediction mechanisms for crowd densities.

(H2): Such trained prediction mechanisms can forecast the density of crowds with acceptable precision in a defined time frame.

Previous works to predict crowd density exists, however most use different approaches than the one presented in this paper. [HLT11, MLHT04] use light-based/visual approaches that can only cover areas that were previously equipped with cameras. [NIK08] use wireless signal strength indicators (RSSI) to estimate the number of devices in an area. However the mentioned approach was only validated in a small-scale experiment. Perhaps most similar to our work is the research conducted in e. g., [WFR+13, WL13]. The authors also use collaboratively crowd-sourced datasets of users' spatial properties to predict crowd densities at a given time and location. Our work deviates in that it uses a different forecasting approach to predict crowd density.

18.2 Approximating Human Density - A Data Collection Approach

In order to investigate patterns of crowd movement and resulting crowd densities we performed a large–scale experiment during an urban sports event in the city of Vienna. Our goal was to observe and track movements and neighbourhoods of spectators in the event area. We developed an event app for Android and iOS that was active on the event day and sent spatial data of willing users to a database. The App was active only on the event day between 7 o'clock and 18 o'clock and was used by 6698 people (android: 5192, iOS: 1506) of which roughly 5230 agreed to participate in our study.

18.2.1 Data

In order to collect spatial data from as many persons as possible it is important to choose sensor technologies that are commonly used. We decided to employ GPS sensors, which are widely distributed in mobile platforms and can be easily used to log the position of devices with fair accuracy. The inherent restriction is that GPS cannot be easily used in indoor environments. In order to detect additional devices in the users surroundings, we further used proximity-based sensor technologies to capture the user's neighbourhoods. Bluetooth, NFC, RFID and WiFi have been used in the past – see for example [MBGH07] – however all technologies suffer from various drawbacks such as short interaction range, limited availability or sparse scan points. Since we did not actually transmit information but simply performed presence detection, we decided on using Bluetooth which is very widely used and can detect other Bluetooth devices in a fitting distance with acceptable speed and accuracy. We collected GPS locations in an interval of 1 minute and scanned for Bluetooth devices in an interval of 2 minutes. To improve the quality of the collected movement traces we applied filtering mechanisms to the GPS data thereby removing unrealistic, short or problematic traces. The resulting GPS dataset contains 1503 user traces collected over 11 hours with a sampling rate of one minute. On average, each trace contains 128 samples with a duration of 265 minutes whereby the median of the interval between samples is 1 minute. The Bluetooth data was collected by 3352 active users that detected 38816 unique ids and generated 1092694 proximity relations overall. In Su et al. [SCM+06] the authors substantiate that assuming bidirectional communication links between moving Bluetooth devices in urban environments is feasible.

18.2.2 Similar spatial data-based models

Many similar datasets for spatial data have been collected in previous works. A rather accurate dataset of position traces that represents the movement of taxis in an urban environment was presented in 2003 [JHP+03]. Roughly 500 cars in the city of San Francisco were tracked in a time frame of 30 days. Samples were created in 10 second intervals. In 2009 Lee et al. investigated the properties of human walk patterns [LHK+09]. To this end 8-32 users were tracked in diverse locations such as universities, cities or amusement parks. Locations were collected once every 10 seconds. The durations of the collected traces ranged from a few hours to an entire week. A location-based dataset featuring more users was created by Microsoft Research [ZLC+08]. 128 Chinese volunteers were observed for more than five years with samples being taken every 1-5 seconds or 5-10 meters. Scientific works that presents datasets of Bluetooth relations – collected by commercially available, non-customized sensors – have been published for almost 10 years. One of the first was using the "Serendipity" framework [EP05]. 100 students were using this software for 9 months – mostly on a university campus – and detected 2798 unique Bluetooth devices. Su et al. describe a "war-driving"-based approach to capture the Bluetooth-landscape of Toronto [SCM+06]. To this end, 5 devices were used in 3 different locations in the city, with the result of identifying 296 unique devices. Pietiläinen and Diot presented a Bluetooth dataset

that was collected in the frame of the SIGCOMM 2009 conference [PD12]. 76 active devices were used in a 2 day period identifying close to 16000 unique ids. The sampling rate for the acquisition of this data was 2 minutes. Another dataset describes the Bluetooth landscape of an urban environment [NMS07]. 3 static and 9 mobile devices discovered over 10,673 unique devices and recorded over 362,599 unidirectional relations in a 4 month period. Compared to the mentioned related datasets the contribution of the presented data lies in the very large number of users that move with common purpose in the same environment at the same time.

18.2.3 Observed crowd densities

For the purposes of this research, we created a unified density map of the event area using all the collected GPS and Bluetooth data. We combined the Bluetooth and GPS information using the knowledge that every time a device detected a Bluetooth neighbourhood, the detected device was physically close – which we defined as within a circular area of 30 meters based on standard Bluetooth radio transmission ranges – to the source device. Based on this assumption we placed the detected Bluetooth devices at approximate locations around the source device. We further segmented the event area into squares of 100x100 meters and computed the user densities for all of these segments. This was achieved by aggregating all GPS and Bluetooth information for 2 minute time intervals, computing approximate locations for all unique identifiers in this interval, and counting the number of devices located in each segment. We aggregated data into 2 minute intervals because of the sampling interval of the Bluetooth sensors. The imprecision that was introduced by the approximation of the devices' locations (about 30m) is somewhat mitigated by the comparably coarser segmentation of the city (100m). Figure 18.1 visualizes the described segmentation and computed densities for three time intervals in the morning, midday and evening of the event day.

In the most crowded area of the event – the city center – we measured peak densities of 397 persons per segment (i. e., 10000 square meters). All computed densities were generated using roughly 38000 unique devices (see Section 18.2.1), however not all devices were active during the entire day. Based on estimations of the number of event spectators provided by the organizers around 300000 people observed the event. As such we project roughly 0.31 persons per square meter in the city center during the time intervals in which most of these people were located in the same area.

18.3 Towards the prediction of crowd densities

Tools that are able to predict crowd densities for a given environment and context – such as an urban sports event – give emergency forces and organizers a visual means to detect and further preemptively avoid dangerous situations that could lead to catastrophes [JBH+12] such as the ones mentioned in the introduction of this paper. Towards this end we built a system that predicts the density of crowds in the near future based on observed movement patterns from the past.

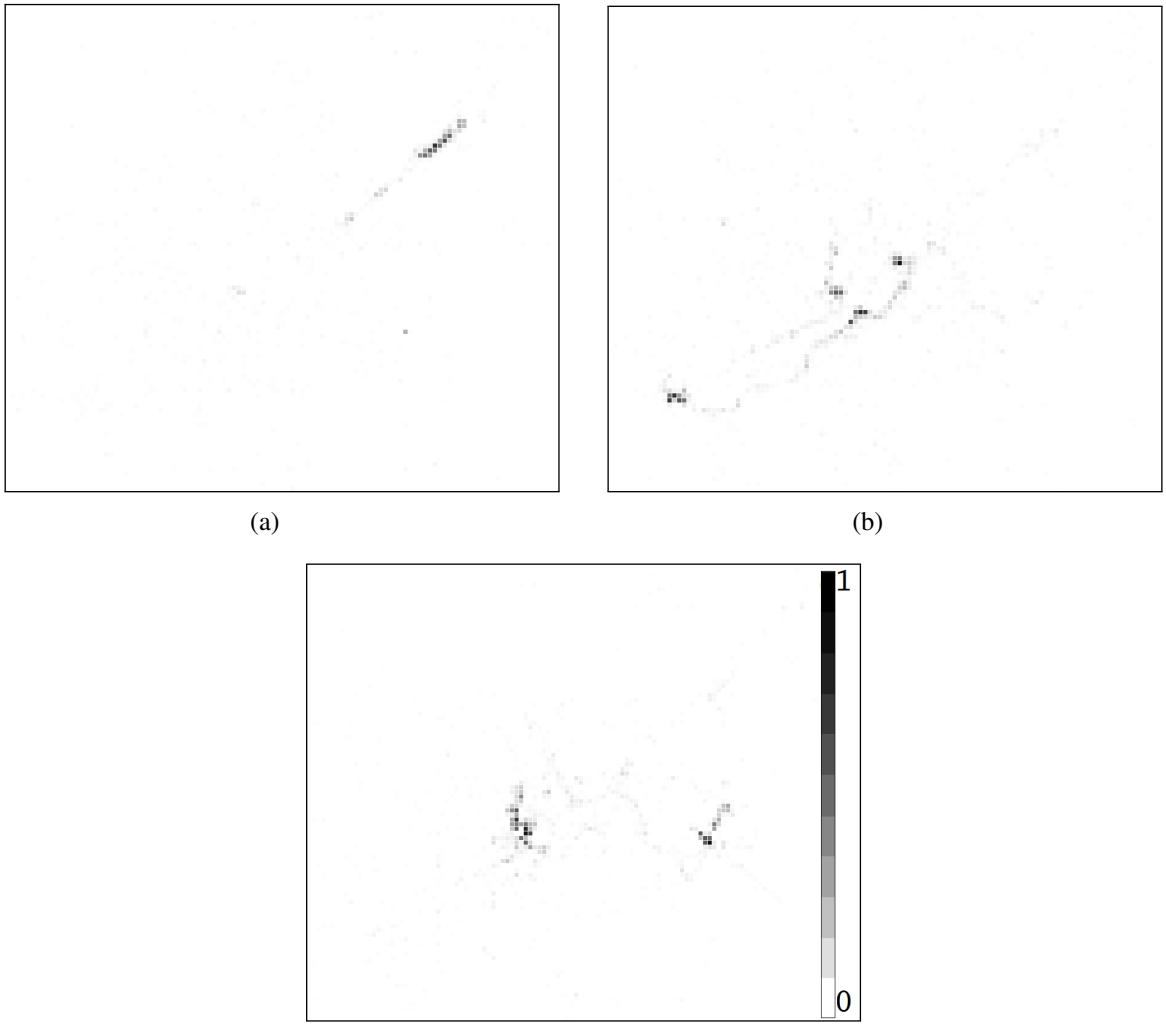

(a) (b)

(c)

Figure 18.1: Segmentation and crowd densities for (a) the morning, (b) midday and (c) the evening of the event day. In this visualization segment densities above 150 spectators per segment were clipped to 150. All values were subsequently divided by 150 and mapped to gray scale.

We approach this problem both by attempting to model the development of individual segment densities directly as well as by predicting the movements of people between segments using graph theory and network flow analysis.

18.3.1 A human crowd density flow network

We build the graph of the city by interpreting segments as vertices and connecting them with two unidirectional edges between neighbouring segments thus creating a checkboard-like network structure. We abstract the movements of people in such a way that they can only travel along the created edges. In order to determine the flow of people between segments during the event, we computed the difference of the segments' densities between consecutive time steps and used the resulting "δ-densities" as starting values for the traditional minimum-cost flow/transportation problem [Mon81]. Segments with positive δ-densities were hereby treated as flow sources, those with 0 δ-density as transport nodes and segments with negative δ-densities as sinks. All edges between vertices had unlimited capacity and cost 1 associated with them. However due to the nature of the data used to generate this flow network (mobile devices that were turned on or off at random), the resulting flow network cannot be guaranteed to be balanced. As such we created a dummy node that provided/consumed the difference in source/sink flow. The dummy node was connected via network edges to all segments, however to not corrupt the results, all flows over these edges were associated with a cost of 100000 (i. e., so high that such flows are only computed once all normal sources/sinks have been satisfied). Figure 18.2 visualizes the computed network flows for three time intervals in the morning, midday and evening of the event day – captured at similar times as in Figure 18.1. We computed these flows using an established algorithm [BEN$^+$04] to solve linear programming challenges such as the minimum-cost flow problem.

As such we created a time series of segment density distributions in 2 minute time intervals – from here on referenced as $d_{0,2}, d_{2,4}, ...$ – and a sequence of flow networks that model how these distributions convey into one another – referenced as $f_{0,2}, f_{2,4}, ...$ –, whereby $f_{0,2}$ describes the flow that runs between $d_{0,2}$ and $d_{2,4}$.

18.3.2 Prediction Approaches

In order to forecast the development of human densities in the entire tessellated event area we attempted to predict both the flow of humans between individual segments based on their δ-densities as well as the segment densities directly. To this end we conducted experiments using various classification methods. Our attempts use the collected and previously presented dataset for training. The created models were used to predict densities for various time intervals. Performance evaluation is performed in relation to the other approaches as well as in relation to the actual measured densities during the event day.

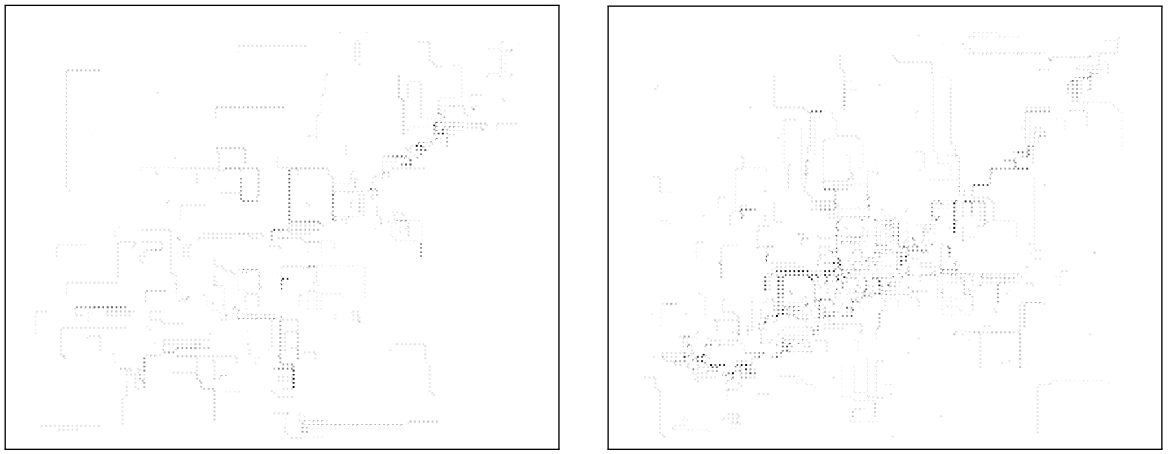

(a)

(b)

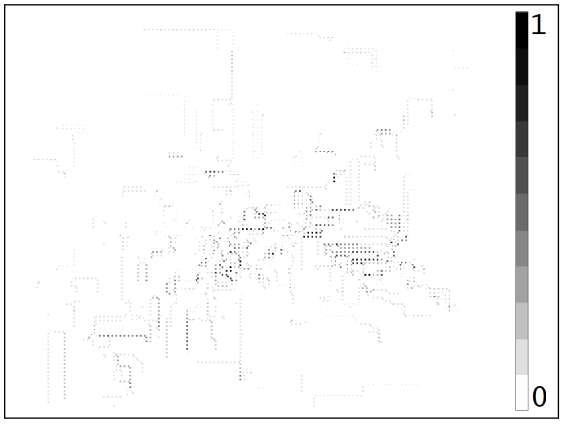

(c)

Figure 18.2: Visualization of the flows between segments for (a) the morning, (b) midday and (c) the evening of the event day. In this visualization flows above 8 were clipped to 8. All values were subsequently divided by 8 and mapped to grayscale.

A first exploratory approach:

To create a baseline comparison using a simple prediction method and in order to investigate the forecasting possibilities using the given dataset, we first conducted experiments using an exploratory forecasting approach. For this purpose we assume that flows between segments are only influenced by the densities of the source and destination segment. Our goal is to compute a prediction function that takes these two densities as input and returns the most likely flow between these segments. To this end we parse the computed values of all arcs of all the generated flow networks and combine them with the associated densities of the source and destination segments for each arc. More precisely, each density pair $(d_{source}, d_{destination})$ taken from $d_{x,x+2}$ is mapped to the respective flow taken from $f_{x,x+2}$. If such a mapping already exists, we add the new reading to the list of previously observed flow values. Further we compute the mean over the entire list of flow values and use it as the predicted flow for the respective density pair in the future. Subsequently we cluster the generated function in order to generate mappings for values that were not part of the training data. We do so by computing the nearest density pair for which we have observed a flow value – using the root mean squared distance metric – and attributing the observed flow to the new density pair. Figure 18.3 visualizes the clustered flow prediction function for density pairs plotting the source density along the y-axis, the destination density along the x-axis and the resulting flow mapped to gray scale.

We can subsequently use this static function to forecast the development of an arbitrary density distribution over multiple time steps. The results of such an analysis are given in Section 18.3.3.

Time series forecasting:

A deeper analysis on the dataset has been performed by means of time series forecasting. We divided our dataset in multiple smaller datasets, each one considering the data of a single segment and its context. For each single segment, in fact, we can include different contents:

1. density of the single segment, and nothing else;
2. density of the segment and densities of the surrounding segments;
3. density of the segment, densities of the surrounding segments and flows on the edges connecting them, as previously estimated in Section 18.3.1.

Moreover, the density of the segments can be measured in two different ways:

1. using the actual absolute density value;
2. using the δ-density value as described in Section 18.3.1;

Consequently, we obtained six different data sets, summarized in Table 18.1. This data partitioning is meant to answer the following questions.

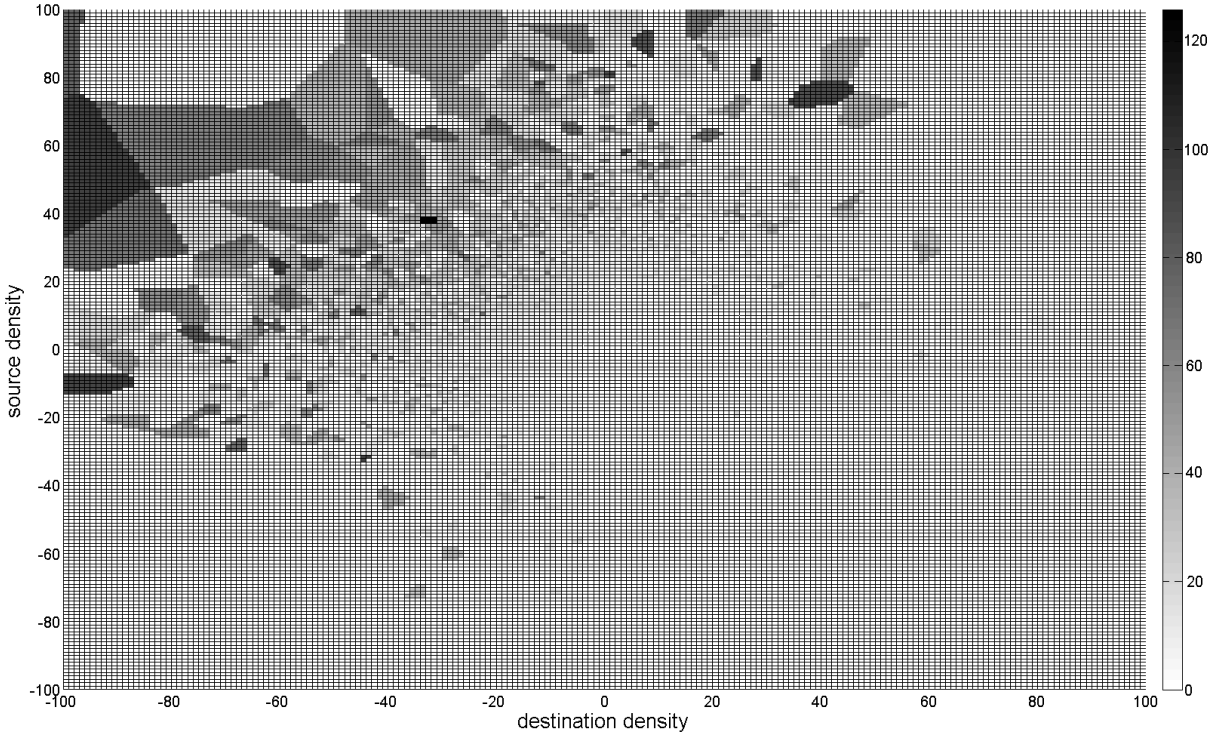

Figure 18.3: A flow prediction function, mapping pairs of segment densities to the expected flow between them. We plot the density of the source segment along the y-axis, the density of the destination segment along the x-axis and map the flow between them to grayscale in the range of 0 to 120.

	Actual density	**δ-density**
Single Segment	SA	SD
Segment and surroundings	NA	ND
Segment, surroundings and flows	FA	FD

Table 18.1: Summary of the datasets used for performing time series forecasting. Each combination is assigned a short name, that will be used from now on.

1. Can a forecaster built upon the information of a single segment provide reliable forecasts, or is the information about its surroundings necessary?

2. How useful is the flow estimation performed in Section 18.3.1? Can it really help in building efficient forecasters?

3. Which measure is more suitable for training a forecaster, the actual density value or the δ-density?

In order to perform time series forecasting, we relied on the tools provided by the well-known WEKA software [HFH+09]. Since version 3.7.3, in fact, such software offers the possibility to perform time series forecasting by means of a dedicated plug-in, that also embeds tools whose goal is to measure how good the forecaster can perform. In this case, no other data set but the one already presented could be used as test set. Consequently, we performed an evaluation of the forecasters using the training data, namely, once the forecaster has been trained, it is then applied to make a forecast at each point in time by stepping through the data. These predictions are collected by the evaluator provided by WEKA, and are compared to the actual data in order to estimate the error that would have been obtained if such a forecaster had been used. We identified three different learning algorithms suitable for our purpose and already available in the tool.

- Linear regression. According to the WEKA documentation, this strategy uses Aikake criterion [Aka74] for model selection, and is able to deal with weighted instances.

- Gaussian processes [Mac98], which has been proven to be flexible and practical to work with in machine learning [Ras04].

- Support Vector Machines [CV95]. WEKA provides a built in utility which implements the sequential minimal optimization algorithm [SS04, KSBM01] for training a support vector regression.

Considering the datasets introduced above, we recap all the combinations of datasets and algorithms in Table 18.2.

	SA	**FA**	**NA**	**SD**	**FD**	**ND**
Linear Regression	SA-l	FA-l	NA-l	SD-l	FD-l	ND-l
Gaussian Processes	SA-g	FA-g	NA-g	SD-g	FD-g	ND-g
Support Vector Machines	SA-s	FA-s	NA-s	SD-s	FD-s	ND-s

Table 18.2: Summary of all the combinations of algorithms and datasets used to perform time series forecasting

18.3.3 Results

Measures

In order to investigate the precision and usefulness of the approaches described in Section 18.3.2 we computed and evaluated their respective prediction errors. For the purpose of this contribution, in which we compare the predictive capabilities of multiple approaches using only a single time series, we use the RMSD (root-mean-squared deviation) metric for comparisons. This gives the inherent benefit that the results are on the same scale as the original time series. More precisely, for a predicted density distribution $fd_{x,x+2}$ and the actual, observed density distribution $d_{x,x+2}$ we compute the prediction error as

$$E = \sqrt{\frac{\sum_{i=1}^{n}(fd_{x,x+2}(i) - d_{x,x+2}(i))^2}{n}}$$

where n equals the number of segments.

Time series forecasting results

Initially, we compared all the time series forecasting algorithms. Figure 18.4 provides a good insight, by showing in six separated charts the error produced by all the combinations listed in Table 18.2. It is immediately clear that the support vector machines is the most promising approach. In fact, it obtains the best results with all the datasets, except with SD and FD, where the gaussian processes algorithm performs better. Due to the non-linear nature of the dataset, the linear regression algorithm is consistently performing worst. It seems clear that without data transformation it is inadequate when it comes to forecasting the evolution of this complex scenario. The last interesting fact is that all those algorithms tend to perform acceptably well for a few time steps, then their performance suddenly diverges making them completely unreliable.

We decided to focus on the support vector machines algorithm in the following analysis, since it has been proven to be the most promising one. We want to understand how useful is the information about the neighbourhood of a segment, and which is the positive impact, if any, of the information about the flow esteem we extracted in Section 18.3.1. Figure 18.5 shows how the information about the segment itself is not descriptive enough, and when used alone to train the forecaster the predictions are significant only for the immediate future. The accuracy improves a lot when the the values of the neighbourhood are also considered in the prediction. Adding the information about the flows, instead, produced some unexpected results: it does not improve the behavior in case of analysis using the absolute density, while it worsens the forecaster performance when the δ-density is used. This suggests that estimating the people flow by means of solving a minimum cost flow problem could be an oversimplification, and together with the fact that such data increases the number of inputs for the regression function, may lead to consistent errors.

The comparison between absolute density and δ-density shows different behavior depending on the data set: with the single node, the forecaster trained with the absolute values seems to

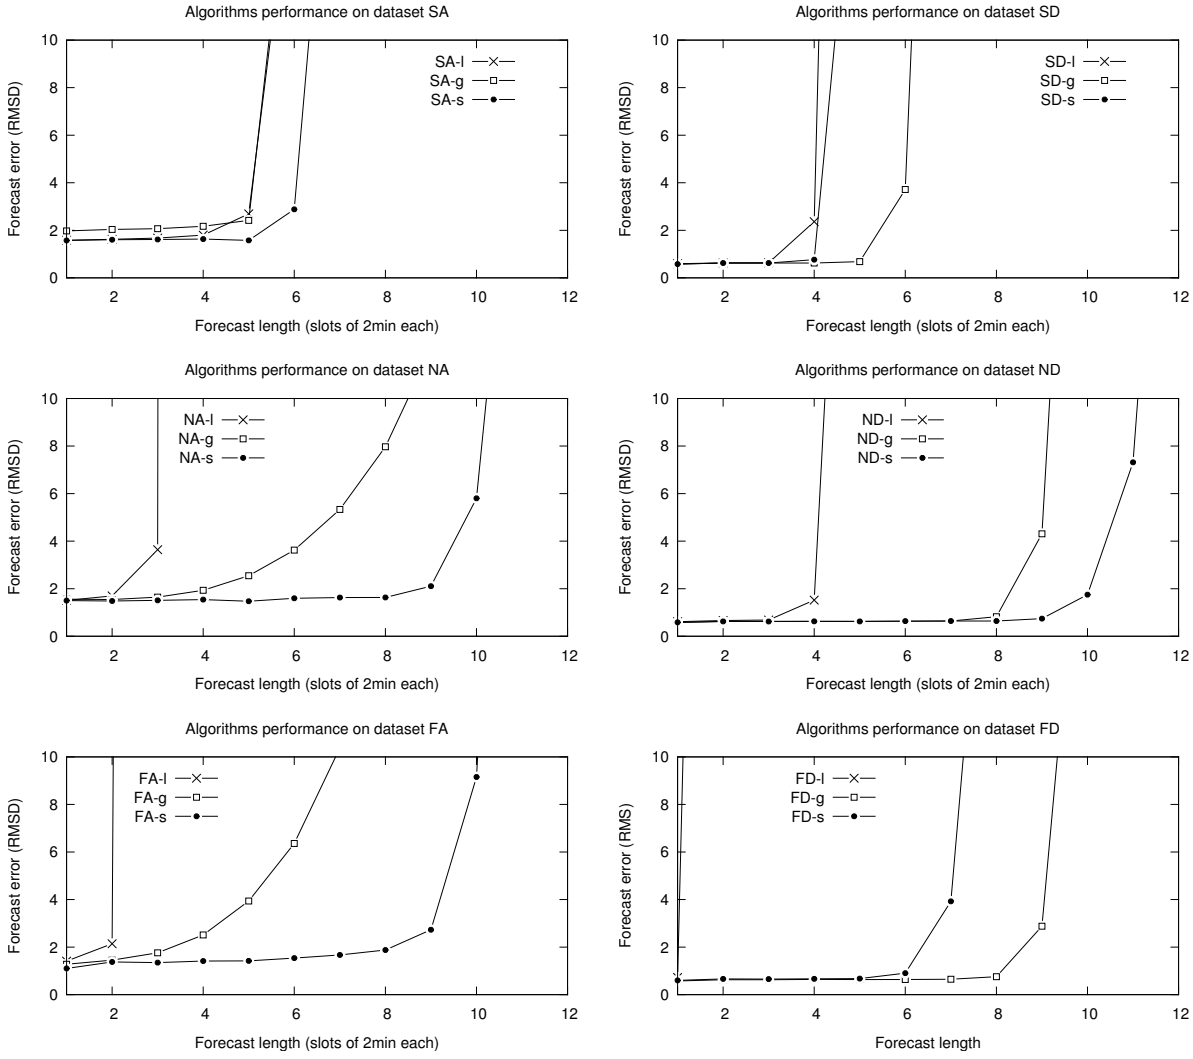

Figure 18.4: Performance of the three algorithms with all the available datasets

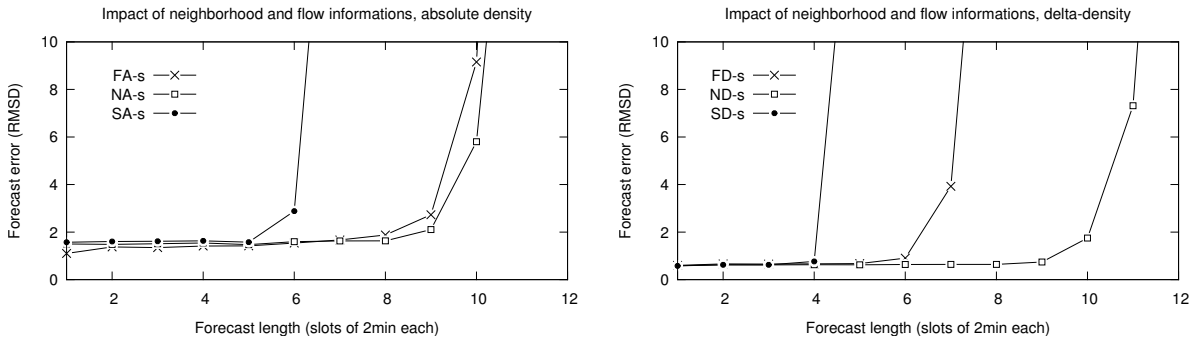

Figure 18.5: Impact of the data about neighbouring segments and flows on the performance of the support vector machines algorithm

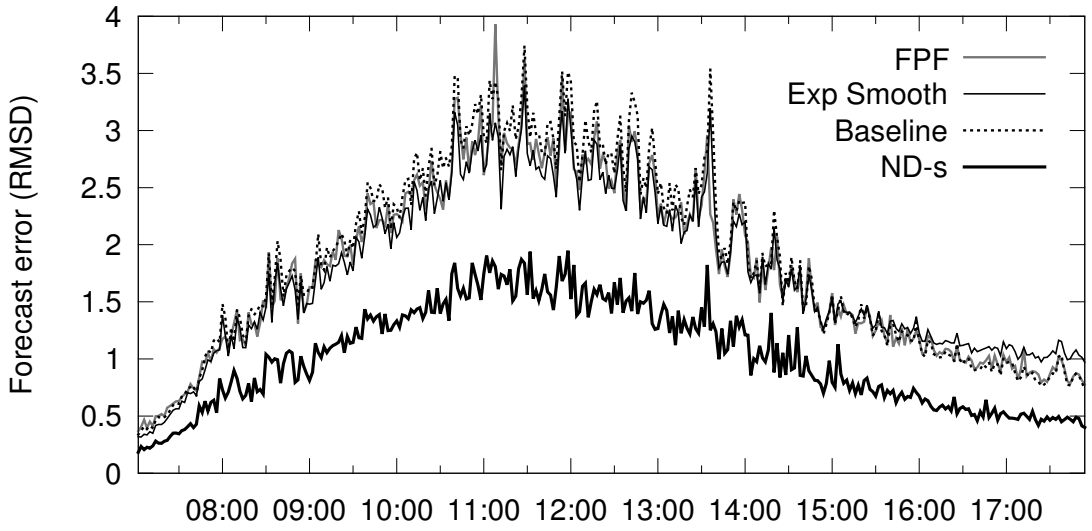

Figure 18.6: Accuracy of various approaches for a forecasting depth of a single step.

perform slightly worse, but this measure is more stable when the forecast window is wider; with informations about estimated flows and neighbourhood, the result is curiously exactly the opposite. Without the flow information, the δ-density performs better. Our hypothesis is that this is due to the fact that δ-density implicitly carries information about the flows when neighbouring nodes are considered. Consequently, it does not benefit from the additional information carried by the our flow estimates.

Comparison of various approaches

Figure 18.6 compares the prediction accuracy for a forecasting depth of a single step. The analysis was performed for the entire collected time series, forecasting the next distribution (i. e., the crowd density in 2 minutes) and comparing the result to the actual crowd density at that

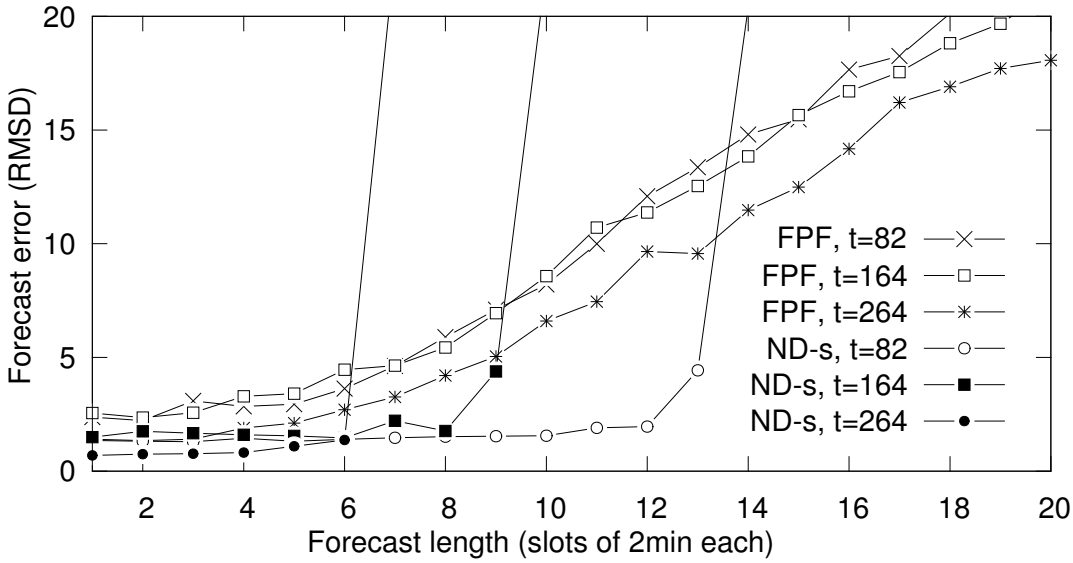

Figure 18.7: Comparison of flow prediction function and support vector machines forecaster for various forecast lengths at different time of the day

time. The figure compares our approaches to a simple baseline algorithm that always predicts the previous density distribution and to the performance of Brown's simple exponential smoothing algorithm [Bro04]. Flow prediction function (FPF), exponential smoothing and baseline perform very similarly, while the ND-s forecaster is able to obtain better results, in particular at midday when the people movement is more intense.

We run the flow prediction function and the forecaster based on support vector machines at the three different times of the day presented in Figure 18.1. We chose to predict up to twenty time steps, and results in Figure 18.7 make clear that the support vector machines algorithm works better for predictions in a short time window, then it completely diverges. On the other hand, the flow prediction function is very stable, but is affected by a bigger error when performing short-term forecast.

Due to the fact that the approaches described in Section 18.3.2 were trained and evaluated on the same dataset, their results are likely affected by training bias. This is a double edged sword: on the one hand these approaches provides good results for future instances of the given urban sports event and similar occurrences, on the other hand if applied to significantly different case studies they may produce less reliable results.

Algorithms performance

We measured the time required to train the forecasters, in order to understand how much computational power would be needed in order to continuously prime them with real time data.

On our test bed[2] the training time for the ND-s was of about two minutes and thirty seconds, thus allowing for a near real time usage. However, the training phase is much longer for datasets in which the flow information is considered. The flow prediction function is faster and perfectly suitable for real time prediction: the training phase only takes around ten seconds.

[2]Intel Core i7 2670QM @ 2.2GHz, 8GB RAM, Sabayon Linux 13.04, Oracle JDK 7u25

19 Case study: Crowd steering at the urban scale

In section, we focus on the problem of crowd steering at the urban scale. The problem is noticeably different from the classic per-user navigation: we want to consider the current position of people to make the user avoid congested areas, thus reducing the overall trip time and improving the security level. Also, this differs from the experiment in Chapter 15 for the complexity of the environment and the number of devices involved. Crowds of people vary with time, and consequently we need a system able to dynamically adapt to such changes. As per our line of research, we obviously want the computation to happen in a completely distributed fashion, with no centralised computing system involved.

19.1 Devices and physical configuration

We suppose users to be equipped with smart devices, also able to communicate with other devices within a certain range with a wireless technology, and must be able to be aware of their current position.

We also suppose the organisers to have spread around the city some "static" network nodes, for instance on public illumination poles, traffic lights, or signals. Such nodes must be equipped with network capabilities similar to those of a Wi-Fi access point. In particular, they must allow nearby mobile devices to connect to all the static nodes within their communication range. There is no upper or lower number or density (devices per square meter) limit to the number of static devices to deploy. The only strict requirement involves both the distance among devices and their communication range: the devices must be placed at a distance and have a communication range large enough such as there is no network segmentation. Besides strict requirements, some other guidelines apply:

- it is preferable to displace static nodes in points where there are people, e.g. on streets and crossings;
- the wider the number of static nodes that can communicate, the more precise will be the results;

- the higher the density of static devices, the more precise will be the results.

The last assumption we make is that devices are able to estimate the time needed for walking towards any of the connected device in normal conditions. This last assumption is again perfectly acceptable considering the current technology, where smart devices are always connected to the Internet and can access public navigation services.

19.2 Distributed crowd steering

Given the structure described above and the mobile smart devices able to communicate with the static nodes, we now outline a possible software solution. The device wanting to be steered publishes a gradient including the information about the wanted destination. The information about destination can be either expressed as physical location (e.g. latitude and longitude) or as a description: the system works in a fashion totally orthogonal to this choice. One or more static devices located near a potential point of interest (POI) receive such gradient, and react becoming sources of a gradient which measures the distances between nodes. This response gradient diffuses on static nodes until it reaches the source, carrying within itself information about the chain of static devices forming the shortest path. The requester will receive a gradient pointing towards the nearest POI, and can navigate step-by-step by reaching in order all the nodes of the path. Although it is not part of this very experiment, the reply could be tuned to be different depending on the request and on the characteristics of the point of interest. For instance, a value representing the "level of matching" between the request and the response could be embedded in the gradient, and used to modify its spatial structure in such a way that POIs more affine with the request are preferred even if farther. Some example of similar usage of gradients is available [SYD+13].

Now that a basic form of steering is in place, we can improve it by adding contextual information. We assume that static nodes are able to detect the number of people in their surroundings. Again, this is not critical even in a real deployment, and can be achieved in numerous ways with a different level of precision, which spans from just keeping track of which mobile nodes are connected (and, as a consequence, within the communication range) to the usage of dedicated sensors and techniques (e.g. cameras and computer vision). Let's call C the perceived number of mobile devices surrounding a static node. We can alter the spatial shape of the gradient by acting on the function that outputs the actual distance between devices. For instance, adopting the notation in Section 3.3.1, we can use $f(n) = \Gamma_n + d(n) + K \cdot C$ where K is a system parameter. What happens? Basically, areas densely populated with mobile devices will appear as more distant, and consequently the requester will be steered towards alternative routes. The whole system works in a totally emergent and distributed way, with no global knowledge of what's going on, no central control involved, and complete distribution of the computation among its components.

A further refinement of these steering mechanisms involve the ability to react to events that will happen in future (given that there is information about that). The patterns required to do

so have already been presented in Chapter 13, and the interested reader can deepen reading [MPV12a].

19.3 Simulation in ALCHEMIST

We decided to focus on the case of mass urban sports events, and in particular on the Vienna City Marathon 2013, an event that every year involves about 40.000 actives and 300.000 spectators. During such event, a smartphone application based on SAPERE [ZCF+11] concepts was deployed, and gathered 1503 high quality GPS traces [APNF13]. We relied on such data set to build simulations demonstrating our concepts of crowd steering at a urban level.

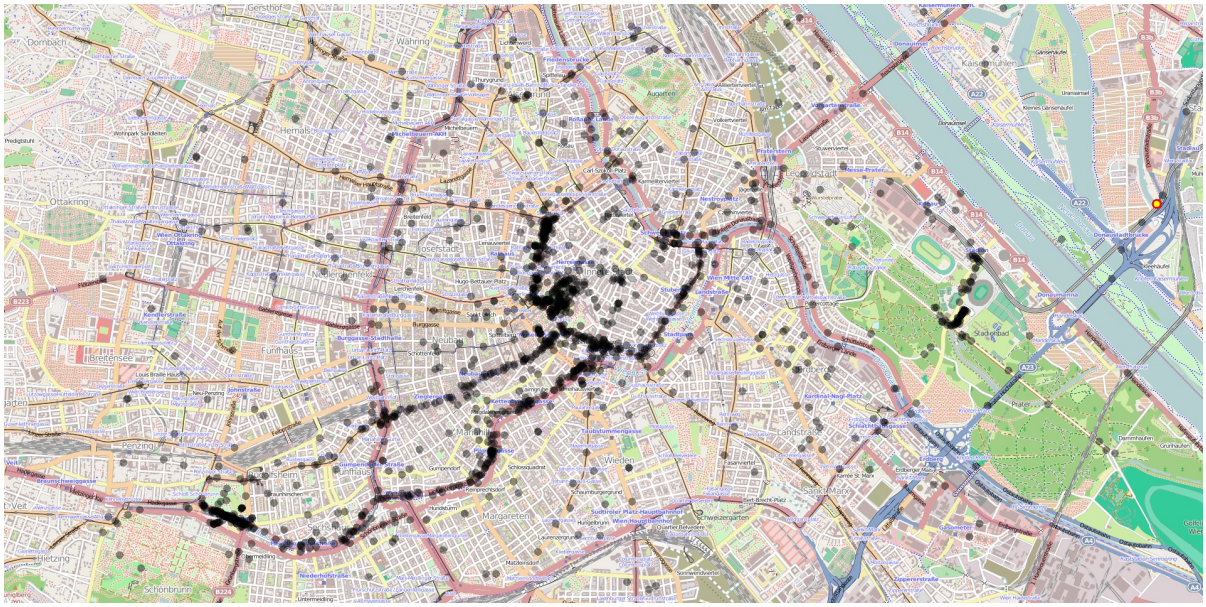

Figure 19.1: A snapshot of the whole city of Vienna as simulated in ALCHEMIST. This snapshot is taken while simulating the city at 10am, each black point corresponds to a GPS trace. The more an area is crowded, the blacker it appears in the image.

As ALCHEMIST environment, we obviously used the map of Vienna. Another rather straightforward choice was to map each device to a node, supposing the users to carry one and one only device which participates to the system. The network of static nodes was built by creating a grid of 572 devices spread around the city and enabling the positioning in the nearest street point. After that, we positioned a node at Burggarten (48.203926,16.365765), representing our POI, and a node at (48.21441,16.35825) representing our navigation system user. Finally, we positioned the nodes of the users which follow the traces. As result, we had in total 2077 simulated nodes, of which 1504 mobile and moving during the actual simulation, 1503 by following the traces using the mixed mode (see Figure 8.3(c)) and one following the steering system.

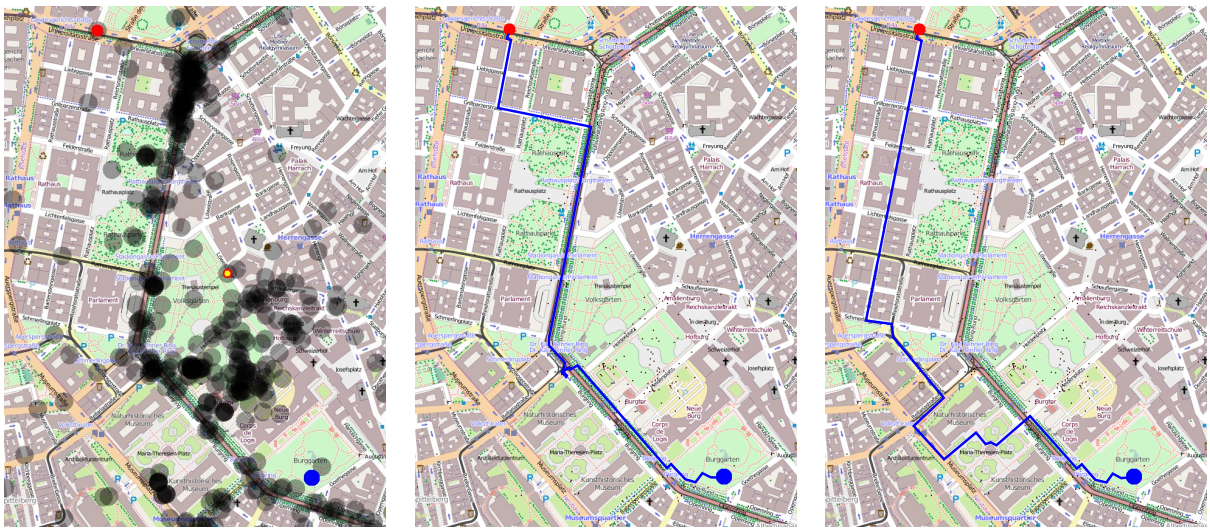

(a) Level of crowding in the interested area

(b) Classic navigation

(c) Crowd-sensitive navigation

Figure 19.2: In these three snapshots a qualitative evaluation of the crowd-sensitive steering system benefits is offered. The requester is painted in red, while the nearest POI is in blue. In Figure 19.2(a), the density of people in the area is shown (at the time at which the red node starts): a more intense black mean a more crowded area. In Figure 19.2(b) and Figure 19.2(c) is possible to see the suggested routes of, respectively, the classic navigation and the crowd-sensitive navigation. The second suggests a longer but much less jammed path.

As linking rule we decided to let the infrastructure nodes connect to every other node within a communication range of 150 meters. We used the same set of reactions for each node in the system, but on the mobile nodes, in which also the reactions needed to program the mobility were included. In Figure 19.1 a graphical output of the simulator is proposed. We relied on the SAPERE meta-model. This way, we were able to simulate a network of programmable tuple spaces.

A complete evaluation of such an emergent environment is all but trivial. The best way to measure the effectiveness of a crowd-sensitive user steering system would probably be the measure of the average walking time on routes generated probabilistically according to the most walked paths. This kind of evaluation has two strong requirements to prove itself scientifically relevant: first, it requires to identify how popular are the routes; second, and most important, it would require a realistic model of pedestrian, including the physical interaction with other people. At the time of writing, we do not have such model, although we are working to find a decent approximation.

As proof of concept, we chose to monitor the path of a single user steered from Universitätstrasse to Burggarten. We chose this path because the shortest path connecting the user to her destination includes a walk in between one of the most crowded areas during the sport event,

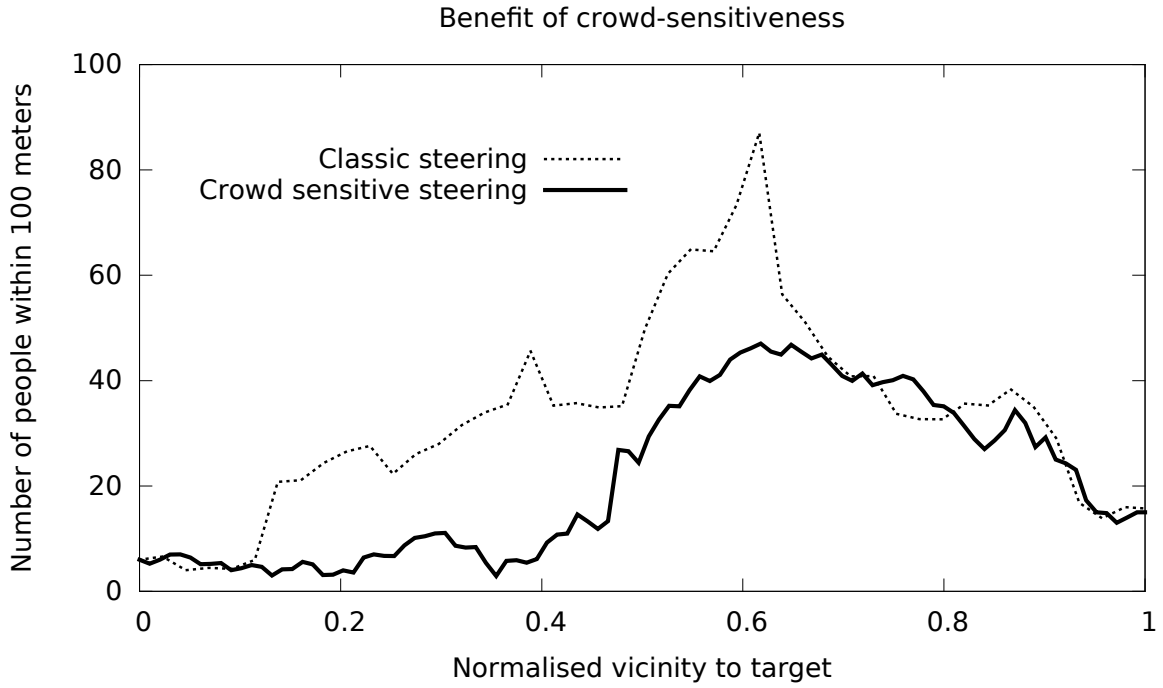

Figure 19.3: This chart shows how the number of users surrounding the user (within 100 meters from her) vary with the proximity to the target. With "normalised distance", we mean that we divided the distance the user still has to walk to reach the target by the total length of the suggested path.

and as such is the perfect test bed for an approach whose goal is to improve the trip time in such conditions.

In Figure 19.2 a qualitative impact is given. It is immediately clear that the suggested path in Figure 19.2(a) is longer but successfully reduces the space walked within crowded areas, with respect to the one in Figure 19.2(b).

The chart in Figure 19.3 confirms the qualitative evaluation. Towards the beginning and the end, the two algorithms have a similar performance, due to the fact that these parts of the walk (as the initial) are common. In the central part of the walk, however, the path suggested by the crowd-sensitive algorithm tends to avoid a considerable amount of jammed areas. We claim that, assuming the traces to be a representative sampling of the actual population of the area and the more crowded areas to be slower to walk (both assumptions sound rather straightforward to the authors), then the crowd sensitive algorithm guarantees an average lower time to destination.

19.4 Final remarks

In this part of the thesis we have focussed on methods and patterns for self-organisation. We contributed with a process algebra for SAPERE, upon which we build as a proof of concept a non-trivial spatial pattern (the channel). We also devised a new pattern devoted to anticipate future events and react accordingly. In order to work correctly, such pattern needs to have information about future events: relying on an existing database of a massive urban sport event, we have studied ways to predict the possible formation of dangerous crowds. Also, we provided a serie of case studies that comprise crowd steering and evacuation, self composition, and semantic resource discovery.

Part IV

Aggregate programming languages

20

Tuple based aggregate programming

In this part, we take a more fundamental approach and try to tackle the issue of engineering computational ecosystems by working at the language level. Working at this level implies identifying proper abstractions and being able to provide a mean for the developer to directly write using those abstractions rather than lower level building blocks. We initially try to extend the existing tuple based coordination toward space and time programming. As a second step, we reason about programming space and time rather than devices, and about what happens when device density changes: can we build a programming abstraction that works regardless the number of devices disseminated in the space? Finally, we introduce Protelis, the last contribution of this thesis. Protelis is an aggregate programming language that inherits the operational semantics of Field Calculus and provides practical tools: a working interpreter and virtual machine, and an integration with ALCHEMIST, that allow the very same code to be tested in the simulator and deployed in real devices.

A common viewpoint in developing coordination models on top of tuple spaces is that tuples are a mechanism to reify and exchange events/data/knowledge that are important to system coordination, and to synchronise the activities of coordinated components in a parallel and/or distributed system. A tuple is, however, a very point-wise abstraction, and applications often need to express relationships articulated across the physical environment through which the computational system is distributed. There is thus a need for coordination models and languages that raise the level of abstraction from the single tuple to a spatial structure of tuples, without forgetting the possibility that such a structure, as well as being distributed, may also be highly dynamic and mobile.

Based on existing works [VCMZ11, MZ09, NFP98, PCBB07], and on an existing trend of studies of concurrency "in space" [CG00, CG10, BMS11, BB06], we introduce a new coordination model and language aiming at further bridging the gap between coordination and the continuum ideas of spatial computing. In our model, we make situated agents interact by injecting into the coordination substrate so-called *space-time activities*, namely, processes that manipulate the space-time configuration of tuples in the network. Such activities are expressed in terms of a process-algebra like language, composing atomic coordination primitives in the style of Linda with additional Proto-derived constructs that deal with spatial and temporal aspects:

1. spreading of an activity in a node's neighbourhood depending on its relative orientation;

2. scheduling an activity at the next computation round of a node;

3. accessing space-time contextual information to link the configuration of tuples with the actual physical space.

20.1 Tuples in space-time

20.1.1 Basic model

Our coordination infrastructure runs on a (possibly very dense and mobile) set of situated computational devices, e.g., located in specific points (i.e. nodes) of the physical environment. Each node hosts software agents (the coordinated components) and a tuple space, and has ability of interaction with nodes in the neighbourhood—where proximity can be seen as a physical or virtual property. Differently from Linda, in which agents interact by atomic actions (inserting, removing and reading tuples in the local tuple space), in $\sigma\tau$-Linda, agents interact by injecting *space-time activities* (activities for short). These activities are processes composing atomic Linda-like actions with additional constructs allowing the activity to diffuse in space (i.e. to other nodes) and in time (i.e. to be delayed). The net effect of an activity is hence to evolve the population of tuples in the network, thereby affecting the distributed structures of tuples that agents use to for global coordination.

In our model, each node undergoes the following *computation round*:

1. it sleeps, remaining frozen;

2. it wakes up, gathers all incoming activities (contained in messages received either from neighbour nodes or from local agents) and executes them;

3. it executes the continuation of the activity executed in previous computation round;

4. it spreads asynchronous messages to neighbourhood;

5. it schedules an activity continuation for next round;

6. then it sleeps again, returning to the beginning of the cycle.

The duration of the computation round is dictated by the underlying infrastructure, and could possibly change over time or from device to device [BBF07]. We only assume that it is long enough for executing steps (2-3-4) above. This particular computational model for a tuple space has many similarities with the platform assumptions of the Proto language [BB06, VBC11, VBU13].

20.1.2 The coordination language

A key role in the proposed coordination model is played by the concept of space-time activities. We here incrementally describe their features by presenting a surface language for their specification.

Primitive actions

We being with three basic Linda actions for manipulating the tuple space: "`out` *tuple*", "`in` *tuple*", "`rd` *tuple*". These respectively insert, remove and read a tuple from the local tuple space. A tuple is a ground first-order term in our model—similarly to [OZ99]. Read and removal specify a template (a term with variables, denoted as literals starting with an upper-case letter) that should be syntactically matched with the retrieved tuple. Read and removal are predicative: they are non-blocking and yield a negative result if no matching tuple is found. To these, we add a fourth primitive, "`eval` *pred*", which evaluates the predicate expression *pred*.

When such actions are defined for an activity injected by an agent, and are executed in the tuple space where the agent is situated in, then a notification result is shipped to the agent—although, following the spirit of Busi et al. [BGZ97], we shall not discuss internal aspects of agent interactions in this thesis.

Protocols

Primitive actions can be sequentially composed in a protocol-like manner. Other standard operators of parallel composition and choice could be orthogonally added, but are not discussed in this thesis for brevity. Additionally, `in`, `rd` and `eval` define branches leading to two different continuations, one for positive and one for negative outcome of the predicative action. Three examples of activities are:

```
out t(1,2,3); out t(a,1+2,b)
in r(X,2,3) ? out r(X,2,3) : out r(0,2,3)
(in r(X,2,3) ? (eval X=1 ? out r(X,2,3) : 0) : 0); out ok
```

The first expression inserts tuple `t(1,2,3)` and then `t(a,3,b)`. Note that tuples are evaluated before being used in actions: evaluation amounts to computing the result of (mathematical) expressions used in a tuple's arguments.

The second expression removes any tuple matching `r(X,2,3)` (variable `X` is bound to the value of first argument, and this substitution propagates through the remainder of the activity). If it succeeds (`?` branch) the tuple is inserted back, otherwise (`:` branch) a new tuple `r(0,2,3)` is inserted.

The third example attempts to remove any tuple matching `r(X,2,3)`. If it succeeds and $X = 1$ then it inserts it back, otherwise it does nothing (`0`). Independently of the outcome of such a removal, tuple `ok` is then inserted. We may also omit the denotation of a ":" branch when it leads to the execution of empty process `0`, writing e.g.

```
(in r(X,2,3) ? eval X=1 ? out r(X,2,3)); out ok
```

in place of the third example above.

20.1.3 Definitions

When desired, one can equip the specification of an activity with *definitions* (which can possibly be recursive), in the style of agent definition in π-calculus. These have the form "*N* `(`x_1`, . . . ,`x_n`)` `is` *activity*", which define *activity* as having name *N* and arguments x_1,...,x_n. For instance, after declaration

```
in-out(T) is (in T; out T)
```

we have for instance that activity

```
in-out(r(1,2,3))
```

behaves just like

```
in r(1,2,3); out r(1,2,3)
```

namely, the tuple is added if it is not already there. Note that since no branches are used for the removal operation, this in turn is equivalent to

```
(in r(1,2,3)?0:0); out r(1,2,3)
```

or, similarly, to

```
in r(1,2,3)?out r(1,2,3):out r(1,2,3)
```

20.1.4 Time

The language provided so far is still point-wise in space and time. We now expand it, beginning by adding construct `next` to situate activities in time. Executing action "`next` *P*" (where *P* is the protocol – also called process – defining an activity), amounts to scheduling *P* for execution at the next computation round of the current node. Special variable `$delay` can be used in *P* and evaluates to the amount of time passed in between the current computation round and the previous one. In fact, even though the time model is discrete at the device level and measured in rounds, devices are supposed to be able to measure how long a round is, and consequently be able to reason about continuous time (even though they sample it discretely). Similarly, variable `$this` can be used to denote the identifier of the node on which it is evaluated. Useful examples of definitions (along with a brief descriptions of them) are then the following ones:

```
% When a tuple matching T is found,
% it is removed and replaced with T2
chg(T,T2) is (in T ? out T2 : next chg(T,T2))

% Inserts a tuple time(T,X), updated as time X passes
rep(T)  is (out time(T,0); next rep2(T))
rep2(T) is (in time(T,Y); out time(T,Y+$delay); next rep2(T))

% Inserts tuple T that is removed after X time units elapse
outt(T,X) is (in-out(T); eval X<=0
                         ? in T
                         : next outt(T,X-$delay))
```

Note that by the use of `next` in conjunction with a recursive definition, `chg` actually declares an activity with a duration in time, which will be stopped only when a tuple matching `T` is eventually found. Concerning the use of `$delay`, we then observe that – in the spirit of a Proto-style space-time computing model – as the average duration of computation rounds tends to zero, activity `rep` tends to define a continuous update of tuple `time(T,X)` as time `X` passes, and similarly, `outt(T,X)` tends to remove tuple `T` precisely as time `X` passed.

20.1.5 Space

To situate activities in space we introduce construct `neigh`. Executing action "`neigh` *P*" amounts to sending a broadcast message containing *P* to all neighbours, which will then execute *P* at their next computation round. Special variable `$distance` is also introduced, which evaluates to the estimated distance between the node that sent the message and the one that received it. Similarly, variable `$orientation` can be used to denote the relative direction from the receiver to the sender (e.g., as a vector of coordinates [CG10]). Some examples are as follows:

```
% Broadcasts tuple T in the neighbourhood
bcast(T) is (neigh out T)

% Broadcasts tuple T in the neighbourood but only within range R
bcastr(T,R) is (neigh (eval $distance<R ? out T))

% Gossips tuple T in the whole network within range R
goss(T,R) is (eval R>=0 ? (in-out(T); neigh goss(T,R-$distance)))
```

Of particular interest is the last definition, which spreads one copy of `T` to all the nodes whose hop-by-hop distance from the source is smaller than `R`. As in the case of time, as devices become

increasingly dense, and their distance tends to zero, the set of devices holding tuple `T` will actually form a continuous sphere with radius `R` around the origin of `goss`.

Note that it is an easy exercise to define processes dealing with both space and time. For instance, one can define a process `gosst` that adds temporal aspects to the `goss` example, such as to make the sphere of tuples created by `goss` all disappear following a timeout. In the sense of spatial computing interpretation [BB06], the definition of `gosst(T,R,TO)` would be the definition of a geometric space-time activity called "sphere of tuple `T` with radius `R` and timeout `TO`"—useful to limit the spatial and temporal extent of some advertised information.

20.1.6 Finally

We conclude by introducing a construct named `finally`, used to simplify the task of structuring the activities executed at a given round. Executing action "`finally` *P*" makes activity *P* executed in the current round, but only when all the others actually completed. A typical use of this construct is to start an aggregation activity for incoming messages only when all of them have been processed, as in the following equivalent specification of gossiping:

```
gossf(T,R) is (eval R>=0 ? (out T; neigh gossf(T,R-$distance);
                            finally clean(T)))
% Cleans multiple copies of T, leaving just one of them
clean(T) is (in T ? (in-out(T); clean(T)))
```

Messages spread by gossiping cause the receiver to execute the `gossf` activity, which inserts tuple `T`, further spreads messages, and finally schedules the `clean(T)` process for a later time. Only when all such messages have been processed in a round (and there are typically more than one) will the set of all `clean` activities be executed. The result of their execution is that only one tuple `T` will remain in the tuple space. The `finally` construct can thus, e.g., serve a similar aggregation and simplification role to the `*-hood` constructs in Proto.

20.2 Core Calculus

In this section we introduce a formalisation of the proposed framework similar in spirit to other works in literature [BGZ97, NFP98, VC09], namely, by a core calculus taking the shape of a process algebra. As in previous examples, a process algebra is a way for us to formalise our target system with rewriting rules, making sure it is well defined and ambiguities are resolved.

20.2.1 Syntax

Let meta-variable σ range over tuple space (or node) identifiers, x over logic variables, τ over real numbers used to model continuous time, and f over function names (each with a given arity, and used either in infix of prefix notation)—as usual we refer to functions with arity 0

as constants. Meta-variable t ranges over terms built applying functions to variables, numbers, identifiers, and constants, and will be written in typetext font. For simplicity, we shorten special variable `$orientation` to ω, and neglect `$distance` since it can be "compiled away" to term $length(\omega)$ where *length* is a function. We let ε range over evaluations (functions) for terms, write t^{ε} for application of ε to term t, and denote $\varepsilon(\sigma,\tau)$ the evaluation that (other than computing mathematical functions) maps `$this` to σ and `$delay` to τ. For instance, we have `a($this,1+$delay)`$^{\varepsilon(\texttt{id23},5.1)} =$ `a(id23,6.1)`. A substitution θ of variables $x_1,\ldots,x_n$ to terms $t_1,\ldots,t_n$ is expressed by notation $\{t_1/x_1,\ldots,t_n/x_n\}$, and is applied to a term t by syntax $t\theta$, e.g., $a(x,1)\{x/2\}$ means $a(2,1)$. We write $mgs(t,t')$ for the most general substitution θ such that $t'\theta = t$—such a notation makes no sense (as in partial functions) if $mgs(t,t') = \bot$, i.e., when t is not an instance of t'.

Given these premises, the core syntax of the model is expressed by the grammar in Figure 20.1 (a). P defines the syntax of a process (or activity): it includes empty process 0, action prefix, predicative actions with branches, and call of a definition. Note we skipped from this syntax the composition operator ";", which can be basically compiled away once we have action prefix "." and branching "? :"—by straightforward equivalences like $0;P \equiv P$, $(\pi?P:Q);R \equiv \pi?(P;R):(Q;R)$ and $(\alpha.P);R \equiv \alpha.(P;R)$. A space S is a composition, by operator " | ", of processes and tuple sets. The topology of a network is modelled by a composition L of connections of kind $\sigma \overset{t}{\rightsquigarrow} \sigma'$, representing proximity of node σ' to σ with orientation vector t—e.g., expressed as term `coord(x,y,z)` or the like. Finally, a system configuration C is a composition, by operator $\otimes$, of nodes $[S]_{\sigma}^{\tau,\tau'}$ (with id σ, space S, current round at time τ and previous one at τ'), topology L, and messages $P \triangleright \sigma$ (with content P and recipient σ).

Figure 20.1 (b) introduces a congruence relation "$\equiv$", stating when two configurations are to be considered syntactically equal, and hence can be used one in place of the other. First line introduces standard multiset-like properties of operators " | " and "$\otimes$". Second line states that scheduling operators can be lifted out of action prefix placed in parallel with the continuation, and can also distribute in parallel processes. Last line states that when a `finally` and `next` actions are in parallel composition, the latter can enter the former: this will in fact leave scheduling policy for Q unchanged.

20.2.2 Operational Semantics

We define operational semantics by transitions $C \xrightarrow{\lambda} C'$, where labels λ can have the syntax described in Figure 20.1 (a). Label "$\cdot$" means a silent action internal to a node σ; "$\sigma!P$" means device σ is broadcasting a message with content P; "$\sigma\tau?P$" means device σ starts a new computation round at (its local) time τ and still needs to gather messages with content P (at the top level it will take the form $\sigma\tau?0$); "$P \triangleright \sigma$" means an agent is injecting process P in the tuple space σ; and "$L : L'$" means (sub)topology L changes to L' to reflect some mobility or failure in the system. Semantic rules are shown in Figure 20.1 (c).

Rule (STR) defines classical structural congruence. Rules (BRO) and (SND) recursively

t	$::=$	$x \mid \sigma \mid \tau \mid f \mid f(t_1,\ldots,t_n)$	Terms
P,Q,R	$::=$	$0 \mid \alpha.P \mid \pi?P:Q \mid D(t_1,\ldots,t_n)$	Process
α	$::=$	$out\ t \mid \Box P$	Action
$\Box$	$::=$	$next \mid neigh \mid finally$	Scheduling operator
π	$::=$	$rd\ t \mid in\ t \mid eval\ t$	Predicative action
T	$::=$	$0 \mid t \mid (T \mid T)$	Tuple set
S	$::=$	$0 \mid T \mid P \mid (S \mid S)$	Space
L	$::=$	$0 \mid \sigma \overset{t}{\rightsquigarrow} \sigma \mid (L \mid L)$	Topology
C,D	$::=$	$0 \mid [S]_\sigma^{\tau,\tau'} \mid P \triangleright \sigma \mid L \mid (C \otimes C)$	Configuration
λ	$::=$	$\cdot \mid \sigma!P \mid \sigma\tau?P \mid P \triangleright \sigma \mid L:L$	Labels

" $\mid$ " and "$\otimes$" are commutative, associative, and absorb 0

$$(\Box P).Q \equiv Q \mid \Box P \qquad \Box(P \mid Q) \equiv (\Box P) \mid (\Box Q) \qquad \Box 0 \equiv 0$$

$$finally\ P \mid next\ Q \equiv finally\ (P \mid next\ Q)$$

(STR) $$\frac{C \equiv C' \qquad C' \xrightarrow{\lambda} D' \qquad D' \equiv D}{C \xrightarrow{\lambda} D}$$

(SND) $$\frac{C \xrightarrow{\sigma!P} C'}{(\sigma \overset{t}{\rightsquigarrow} \sigma') \otimes C \xrightarrow{\sigma!P} C' \otimes (P\{t/\omega\} \triangleright \sigma') \otimes (\sigma \overset{t}{\rightsquigarrow} \sigma')}$$

(BRO) $$\frac{(\sigma \overset{t}{\rightsquigarrow} \sigma') \notin C \qquad P \not\equiv 0}{[S|neigh\ P]_\sigma^{\tau,\tau'} \otimes C \xrightarrow{\sigma!P} C \otimes [S]_\sigma^{\tau,\tau'}}$$

(REC) $$\frac{C \xrightarrow{\sigma\tau?P|Q} C'}{(P \triangleright \sigma) \otimes C \xrightarrow{\sigma\tau?Q} C'}$$

(NEW) $$\frac{P \triangleright \sigma \notin C \qquad \tau_2 > \tau_1}{[T|next\ Q]_\sigma^{\tau_1,\tau_0} \otimes C \xrightarrow{\sigma\tau_2?P} C \otimes [T|P|finally\ Q]_\sigma^{\tau_2,\tau_1}}$$

(FIN) $$\frac{-}{[T|finally\ P]_\sigma^{\tau,\tau'} \otimes C \xrightarrow{\cdot} C \otimes [T|P]_\sigma^{\tau,\tau'}}$$

(RUN) $$\frac{S\langle P\rangle \xrightarrow{\varepsilon(\sigma,\tau-\tau')} S'\langle P'\rangle}{[S|P]_\sigma^{\tau,\tau'} \otimes C \xrightarrow{\cdot} C \otimes [S'|P']_\sigma^{\tau,\tau'}}$$

(MOV) $$\frac{-}{L \otimes C \xrightarrow{L:L'} C \otimes L'}$$ (AGN) $$\frac{-}{C \xrightarrow{P\triangleright\sigma} C \otimes (P \triangleright \sigma)}$$

(OUT)	$S\langle out\ t.P\rangle$	$\overset{\varepsilon}{\hookrightarrow}$	$(S \mid t^\varepsilon)\langle P\rangle$	
(IN1)	$(S \mid t')\langle in\ t?P:Q\rangle$	$\overset{\varepsilon}{\hookrightarrow}$	$S\langle P\theta\rangle$	if $\theta = mgs(t',t^\varepsilon)$
(IN2)	$S\langle in\ t?P:Q\rangle$	$\overset{\varepsilon}{\hookrightarrow}$	$S\langle Q\rangle$	if $\nexists t' \in S$ and $mgs(t',t) \neq \bot$
(RD1)	$(S \mid t')\langle rd\ t?P:Q\rangle$	$\overset{\varepsilon}{\hookrightarrow}$	$(S \mid t')\langle P\theta\rangle$	if $\theta = mgs(t',t^\varepsilon)$
(RD2)	$S\langle rd\ t?P:Q\rangle$	$\overset{\varepsilon}{\hookrightarrow}$	$S\langle Q\rangle$	if $\nexists t' \in S$ and $mgs(t',t) \neq \bot$
(EV1)	$S\langle eval\ t?P:Q\rangle$	$\overset{\varepsilon}{\hookrightarrow}$	$S\langle P\rangle$	if $t^\varepsilon = \texttt{true}$
(EV2)	$S\langle eval\ t?P:Q\rangle$	$\overset{\varepsilon}{\hookrightarrow}$	$S\langle Q\rangle$	if $t^\varepsilon \neq \texttt{true}$
(D)	$S\langle D(t_1,\ldots,t_n)\rangle$	$\overset{\varepsilon}{\hookrightarrow}$	$S\langle P\{t_1^\varepsilon/x_1,\ldots,t_n^\varepsilon/x_n\}\rangle$	if $D(x_1,\ldots,x_n)$ `is` P

Figure 20.1: (a) Grammar, (b) Congruence, (c) Global semantics and (d) Local semantics

handle broadcasting (mostly in line with similar works [SRS10]), namely, create messages for all neighbours as soon as a process P is scheduled for broadcasting. Rule (SND) recursively selects a neighbour σ' at orientation t, and creates a message for it in which orientation variable ω is substituted with t. Rule (BRO) is the fixpoint: when all neighbours have been handled, scheduling action *neigh P* is removed. Note we do not send empty messages.

Similarly, rules (REC) and (NEW) recursively handle the reception of all messages when a new computation round starts. Rule (NEW) states that, given node σ in which Q is the process to execute at the next round, when a new round starts at time τ_2 and with overall incoming messages P, then the new process to start with is "$P|$*finally* Q", since we prescribe messages to be handled before Q as already described in previous section. Also note that this rule updates round times τ_1, τ_0 to τ_2, τ_1, and that it activates only when all incoming messages have been actually handled. Rule (REC) recursively gathers all incoming messages: it takes one with content P and proceeds recursively adding P to the set Q of messages considered so far.

Rule (FIN) handles semantics of *finally P* construct, by simply stating that when this is the only activity in a node, we can simply execute P—note all "finally-scheduled" processes can be gathered together (along with "next-scheduled" ones) because of congruence. Rule (RUN) handles one-step execution of a process, by simply deferring the task to transition relation $\overset{\varepsilon}{\hookrightarrow}$, defined in Figure 20.1 (d)—its rules are quite straightforward, as they correspond to the standard semantics of Linda primitives in their predicative version [BGZ97]. Note that $\overset{\varepsilon}{\hookrightarrow}$ takes the evaluation function to use, initialised in rule (RUN) with the proper value of `$this` and `$delay`. Finally, rule (MOV) addresses topological changes due to mobility or failures, and rule (AGN) models the injection of a process by an agent in the local node.

We conclude stating isolation and progress properties. First property allows one to reason about the execution of an activity into a node without considering its environment. Namely, we have that nodes get affected by the external environment only at the time a new computation round starts (because of reception of messages), otherwise they proceed in isolation possibly just spawning new messages.

Property 1. *If* $C \otimes [S]_\sigma^{\tau_0,\tau_0'} \xrightarrow{\lambda} C' \otimes [S']_\sigma^{\tau,\tau'}$ *with* $S \not\equiv S'$ *then* λ *is either* $\cdot$, $\sigma!P$, *or* $\sigma\tau?P$. *In the former two cases (namely, unless we change computation round),* $\tau_0 = \tau$, $\tau_0' = \tau'$, *and* $C' \equiv C \otimes C_m$ *(where* C_m *is either* 0 *or a broadcast), and moreover, for each* D *we have also* $D \otimes [S]_\sigma^{\tau_0,\tau_0'} \xrightarrow{\lambda} D \otimes [S']_\sigma^{\tau_0,\tau_0'} \otimes C_m$, *i.e., computation is independent of the environment.*

The progress property states instead that when a computation round is completed it is necessarily composed of a `next` scheduling: at that point (NEW) can surely fire for that node, starting a new computation round. This ensures that our computations never get stuck.

Property 2. $C \otimes [S]_\sigma^{\tau,\tau'} \not\dashrightarrow$ *and* $C \otimes [S]_\sigma^{\tau,\tau'} \overset{\sigma!P}{\dashrightarrow}$ *iff* $S \equiv (T \mid \mathit{next}\ P)$. *In that case, we have* $C \otimes [S]_\sigma^{\tau,\tau'} \xrightarrow{\sigma\tau_0?0} C' \otimes [S']_\sigma^{\tau_0,\tau}$ *for any* $\tau_0 > \tau$.

Crowd-aware gradient

```
% creating a gradient spreading tuple T
source(T) is (in-out(source(T)); grad(T,0,$this))

% gradient process for tuple T, at distance D,
% coming from node S
grad(T,D,S) is grad(T,D,S,$this)
grad(T,D,S,This) is (
  rd source(T)
    ? in-out(pre(T,0))
    : in pre(T,N) ? (eval N<D
      ? out pre(T,N)
      : (in target(T,M); out target(T,S); out pre(T,D)))
); finally (in pre(T,N)
  ? (in field(T,M);
    (rd crowd(C)
      ? out field(T,N-1.2*C)
      : out field(T,N));
        rd field(T,V);
        neigh grad(T,V+$distance,This)))
```

Figure 20.2: Definitions for the crowd-aware computational gradient. At each site, if this is the source we consolidate `pre(T,0)`. Otherwise, we replace the `pre` tuple if a smaller distance `D` is found, and `target` tuple is inserted as well. Finally, we take the remaining `pre` tuple, and apply the `crowd` factor: the resulting distance `N` goes into the `field` tuple.

20.3 Case studies

20.3.1 Adaptive Crowd Steering

As a first example we study a specification able to support the case study presented in other works [VPMS12, PMV11], with the goal of showing how $\sigma\tau$-Linda can provide support to easily define complex, distributed and adaptive data structures, and how they can be used in practice in a pervasive computing scenario.

Our reference environment is a bidimensional continuous space made of various rooms connected by strict corridors. Inside rooms and corridors, a dense grid of computational devices (nodes) is set up. Each node hosts its own tuple space, receives coordination activities (programmed using our spatial language) by software agents running in it, interacts with nodes in its proximity, and has a sensor locally injecting a tuple `crowd(CrowdLevel)` where `CrowdLevel` is an estimation of the number of people sensed around. People want to reach a point of interest (POI) by the fastest path, and receives directions suggested by their handheld device and/or by public displays on the walls. It is worth noting that the fastest path does not

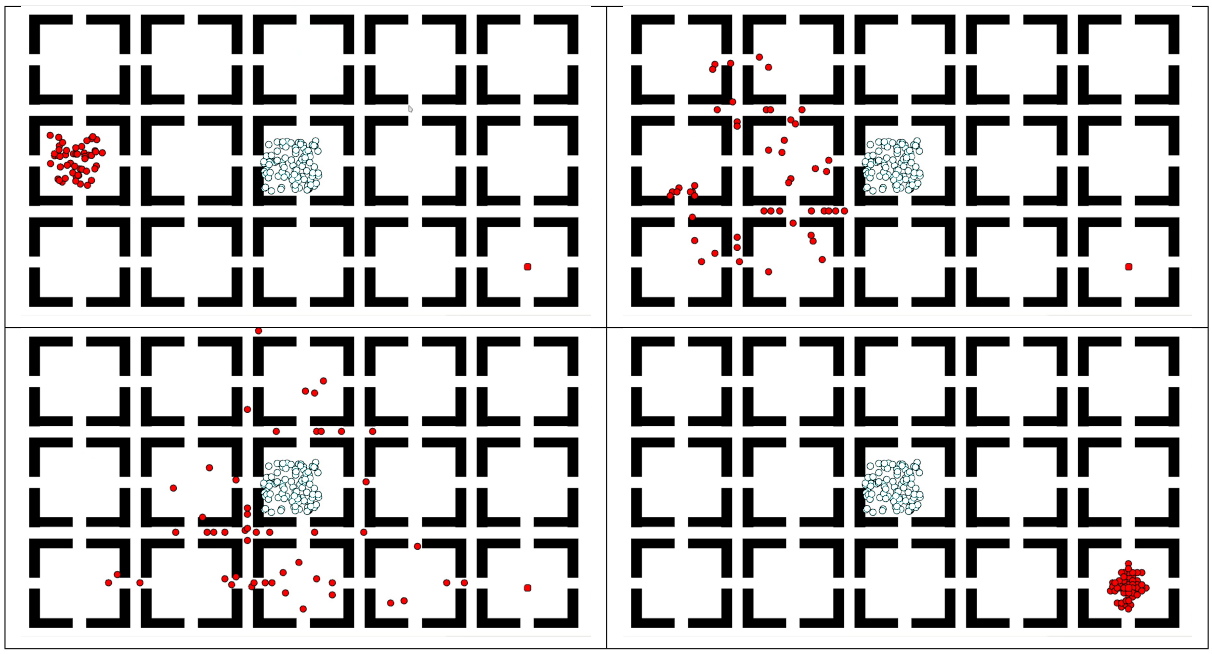

Figure 20.3: Simulation snapshots: the coloured visitors reach its POI avoiding crowd

correspond to the shortest: if everybody followed the same way, in fact, corridors would become crowded. We want the system to be able, relying only on local interactions, to avoid crowded paths, dynamically adapting to any emerging and unforeseen situation. However, we will not implement algorithms to predict future situations, but rather make information about a crowded area spread around such that it becomes a less attractive transiting place to reach a POI.

Our strategy is to build a computational gradient injected by an agent located in the POI. A computational gradient holds in any node the estimated distance to the source by the shortest path [BB06], computed by further spreading and then aggregating at the destination the local estimation of distance. This distributed data structure must take into account also the crowding level, increasing estimated distance where a crowd is forming, and thus deflecting people towards longer but less crowded paths. This strategy can be encoded as in Figure 20.2, where the goal is achieved by maintaining a tuple `target(Poi,Id)` containing the `Id` of the neighbour node where to steer people to following a certain `Poi`. The crowding level influences the local field generation, and is weighted using a constant $K_{crowd} = 1.2$. Values between 1 and 1.5 have been established as good ones after running several simulations: more generally, the higher K_{crowd}, the more sensitive is path computation to the presence of crowd.

We implemented and ran simulations using ALCHEMIST simulator [PMV11], assuming that computation rounds are fired at the same rate for all nodes, and modelling such a rate following the Continuous-Time Markov Chain model. Four screenshots of a simulation run are provided in Figure 20.3, in which we built an environment of fifteen rooms with an underlying grid-like network of infrastructure nodes, an initial configuration with two groups of people, and a POI of interest for the first group which is reachable by a path crossing a crowded area. Note that not

only every visitor reached the POI, but they all bypassed the crowded room (even if it is part of the shortest path, the large amount of people inside makes the whole area rather disadvantageous to walk); additionally, the visitors group is subject to "self-crowding", in that when a group is following a path it forms crowded areas itself (e.g. near doors), hence people behind the group tend to follow different paths. Further simulations we do not describe here for the sake of space show that the above properties hold for a large set of situations, including presence of crowds in different locations and dynamic formation and movement of such crowds during simulation[1].

20.3.2 Tuples in a mobile ad-hoc environment

As a second case study we show a possible extension for Linda standard primitives taking into account both time and spatiality. In particular, our aim is to show how would it be possible in a mobile ad-hoc environment to specify, along with an operation over a tuple, a spatial and temporal horizon of validity: only retrieval operations whose horizon embraces the respective target tuple will actually succeed. We will show an implementation for the spatio-temporal `out` (`stout`) and the spatio-temporal `in` (`stin`) primitives—the easier case of `strd` being a simple variant of `stin`.

The key idea is to make primitive Linda actions actually generate waveform-like space-time data structures, with limited extent in space and dissolving as a timeout expires. Those structures will be responsible to determine the pertinence in space and time of each operation. An example of such a structure is realised by the code shown in Figure 20.4 (top). A `wave` works similarly to the gradient in Figure 20.2, maintaining a `target` tuple reifying the shortest past through a similar specification. A main difference – other than the obvious absence of any crowd management – is the evaluation of the age and distance, which makes the wave disappear whenever and wherever the horizon is reached.

When a `stin` operation requiring retrieval of a tuple template `T` is triggered, it will spawn a messenger activity called `hermes` (with `Op` set to `in`) which will propagate to a matching tuple `T'` following the corresponding `wave` it generated. As soon as the tuple is found, a new `hermes` (with `Op` set to `in_back`) is spawned which will follow the `stin` gradient back. This behaviour can be coded as shown in Figure 20.4 (middle).

Given these two basic bricks, the `stout` and `stin` primitives would be encoded as in Figure 20.4 (bottom). For each, a tuple template, a spatial range and a validity time must be specified. `stout` implementation is concise, because it just needs to manifest itself trough a `wave` and make the tuple available; `stin`, instead, needs also to spawn a `hermes`, whose goal is to retrieve a tuple and move it to the tuple space where the operation was spawned.

These new primitives allow agents to publish/retrieve information flexibly tuning the space-time horizons, relying on lower-level gradients (and routing paths) which adapt to the mobility of the network [Bea09].

[1]The interested reader can download an example clip at: `http://apice.unibo.it/xwiki/bin/download/Publications/Coord2012/museum-small.avi`

Wave-form: a space-time gradient

```
wave(T,Range,Ttl) is wawe(T,Range,0,$this,$this,Ttl, 0)
wave(T,Range,D,Source,Ttl,Age) is
  wave(T,Range,D,Source,$this,Ttl,Age)
wave(T, Range, D, Dest, This, Ttl, Age) is (
  eval (Age>Ttl or D>Range)
    ? (in pre(T,D); in field(T,N))              % disappearing
    : rd source(T)
      ? (in-out(pre(T,0))      % default behaviour in a source
      : in pre(T,N) ? (eval N<D    % choosing minimum distance
        ? out pre(T,N)
        : ( in target(T,_);
          out target(T,Dest);
          out pre(T,D))))
); finally (in pre(T,N)? (              % consolidating target
  in field(T,M);
  out field(T,N);
  rd target(T, Dest);
  next wave(T,Range,D,Dest,This,Ttl,$delay);
  eval Age = 0 ?
    neigh wave(T,Range,N+$distance,This,Ttl,0)))
```

Tuple retrieval

```
hermes(Op, T, This) is
  eval This = $this
    ? (eval Op = in ? ( in T
        ? (rd target(op_in(T), Dest);
          neigh hermes(in_back, T, Dest))
        : (rd target(op_out(T), Dest);
          neigh hermes(in, T, Dest))))
    : (eval Op = in_back ? (in in_request(T)
      ? out(T)
      : (rd target(op_in(T), Dest);
      neigh hermes(in_back, T, Dest))))
```

Space-time Linda operations

```
stout(T,Range,Ttl) is out(T);
  wave(op_out(T), Range, Ttl)

stin(T, Range, Ttl) is out in_request(T);
  wave(op_in(T), Range, Ttl);
  hermes(in, T, $this)
```

Figure 20.4: Definitions of $\sigma\tau$Linda operations.

21

Scale independent computations

Spatially, pervasive networks will tend to be highly variable in their density across space and time, and this represents a key problem. For example, in a smart mobility scenario we can expect that certain areas (e.g., city centers) will host a high number of cars and a very dense deployment of traffic sensors and digital signs, while others (e.g., highways or country roads) will be very sparse. Likewise, the density of devices is likely to vary greatly in time, both as increasing numbers of devices are deployed and as events (e.g., sporting events, festivals) cause large fluctuations in the number of people and amount of accompanying devices in particular areas.

We contend that the development and maintenance of distributed systems in the face of such fluctuations will only be tractable if they make use of techniques that ensure "by construction" that computations are inherently resilient against such forms of variability. Toward that end, the first contribution of this chapter (as an extension of the work of Beal et al.[BVD14]), is the identification of a notion of *consistency* of a space-time computation, stating that the outcome of a computation converges as the density of computing devices increases toward infinity. This is further refined to *eventual consistency*, which focuses only on the stable outcome of a computation (e.g., self-stabilizing patterns as studied by Viroli and Damiani[VD14]).

Second, we identify a set of sufficient conditions for this property to hold, expressed as constructs defining a general purpose spatial language [BDU$^+$13] useful for a wide variety of applications. The language is based on the notation of *computational field*, and it is defined as a restriction of the universal *Field Calculus* introduced in Section 3.5.2 to a subset of expressions that generate only eventually consistent fields, and features a new construct, "gradient-following path integral" (`GPI`), that generalizes a number of frequently used self-organization patterns. We show how `GPI` can be used to implement well-known self-organization patterns such as distance calculation, broadcast, and path forecasting [MVFM$^+$13, VD14], (which are hence also proven to be eventually consistent). Finally, these patterns are applied in several situated network scenarios, and simulations confirm the predictions of eventual consistency.

Symbol	Definition
$\bar{x}$	Closure of a set; sequence of a token
M	A Riemannian space-time manifold
m	A point ("event") in a manifold M
c	Bound on the speed of information propagation
$T^+(m)$	Time-like future: events reachable from m slower than c
$T^-(m)$	Time-like history: events that can reach m slower than c
d	Computational device: a time-like curve in a manifold M
S_M	A space-like cross-section of a manifold M
D	Discrete subset of a manifold, e.g., events in execution of an algorithm on physical devices.
$\mathbb{V}$	Set of all possible data values
f	Field: a function $f : M \to \mathbb{V}$ assigning values to all points in the manifold M.
$\mathbb{F}_\mathbb{V}$	Space of all fields with range contained in $\mathbb{V}$
e	"Evaluation environment" input field
C	Computation: a higher-order function mapping fields to fields: $C : \mathbb{F}_\mathbb{V} \to \mathbb{F}_\mathbb{V}$
$\mathcal{N}(m)$	Communication neighbourhood of m
r	Range within which all points are neighbours
$I^\infty(m)$	Intersection $\mathcal{N}(m) \cap \overline{T^-(m)}$, from which information is directly available to m

Figure 21.1: Table summarizing key symbols that will be used in the following discussion

21.1 Space-Time Computation and Discrete Approximation

This section presents a model of computation over continuous space-time [Bea10, BVD14, BUB13], and defines properties linking it with approximate implementation on a wireless network. The continuous abstraction allows specification of a physically situated computation without consideration of the particulars of the network that will be used to implement it. This in turn provides certain types of scalability and resilience: any computation that is *eventually consistent* (as defined below) produces results that are not sensitive to the precise locations of devices and can only improve (asymptotically) as the number of devices in the network increases. In addition, this abstraction supports higher-level specification of distributed algorithms, since many are natural to describe in geometric terms, e.g., distances, regions, information flow.

As many readers may be unfamiliar with prior results in this area, we begin with a review of the manifold model of space-time adapted from physics [Bea05, Bea10], which then forms a basis for space-time computation [Bea10, BVD14, BUB13]. We then use these concepts to define *eventually consistent* programs and examine the utility of this definition.

For additional assistance to the reader, we provide a table of the symbols (Figure 21.1).

21.1.1 Modelling Networks as Space-Time Manifolds

Many physically situated networks perform computations that can be naturally described in terms of the physical space through which the devices comprising the network are distributed. As observed in literature [Bea05, Bea10], such a network can be viewed as a discrete approximation of the continuous physical space and programmed accordingly.

Under the continuous model developed in those papers, computation takes place on a Riemannian manifold M with both space and time dimensions. A manifold is a space that is locally Euclidean, but may have more complex structure over a longer range. The reader may interpret such manifold as a Euclidean space "extended" with support for "holes" and for arbitrary distance metrics, and where the shortest distance between two points may not be unique. Actually, the only property we must strictly require for a manifold in this work is be paracompactness. However, Riemannian manifolds, which are a subset of paracompact manifolds, also support familiar geometric constructs like angles, lengths, curvature, integrals and derivatives. This allows communication and mobility constraints to be embedded in manifold's geometry, measuring distance through the manifold rather than using absolute (e.g., latitude/longitude) coordinates. For example, the walkable spaces of a building (e.g., Figure 21.2) form a Riemannian manifold in which the shortest distance between locations goes through doors and hallways, rather than through the walls.

Communication is modelled as a bound c on the speed at which information can propagate; given this bound, concepts and terminology from physics well describe the space-time relations of a manifold [TW92]. A point $m \in M$ is termed an "event," denoting its interest as both a spatial and temporal location (see examples in Figure 21.3), and the manifold may be partitioned with respect to m in space and time:

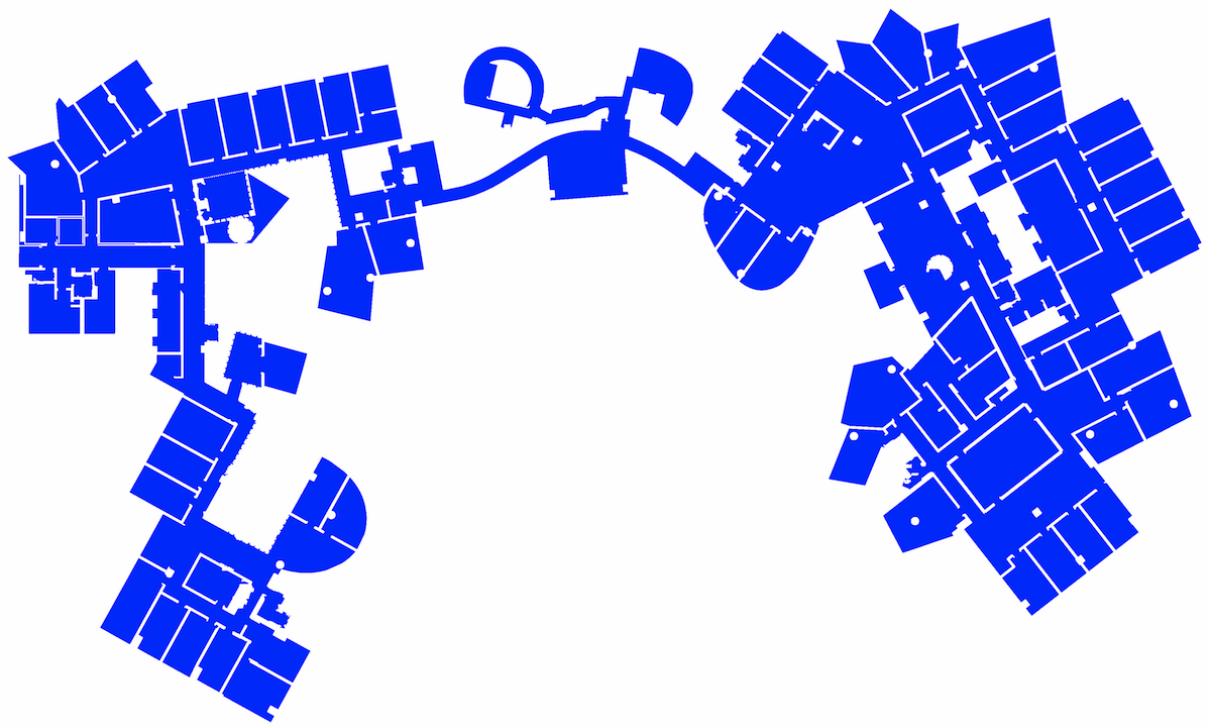

Figure 21.2: The spaces where situated computations take place are often quite complicated, e.g. the MIT Stata Center floor plan shown above. With a manifold representation, distance and other geometric relations conform to the space, e.g., shortest paths follow hallways rather than passing through walls.

- Events that information can go to or from exactly at c are *simultaneous* with m.
- Events that can be reached from m moving slower than c are in its *time-like future* $T^+(m)$.
- Events whose information reach m moving slower than c are in its *time-like history* $T^-(m)$.
- All other events, which cannot share information because it would need to move faster than c, have *space-like separation* from m and no natural order.

A *device* d can be any one-dimensional curve in the manifold with purely time-like relations between the points[1]. The mathematical analysis in this section also makes the simplifying assumptions that the manifold is finite in diameter and that devices do not move[2]. Finally, a *spatial section* is any subspace S_M whose complement is its time-like future and history, i.e., $M - S_M = T^+(S_M) \cup T^-(S_M)$, implying all points in S_M have purely space-like separation. Spatial sections are essentially snapshots in time, without attempting to enforce a global notion of time.

21.1.2 Computations on Continuous Space and Time

Typical formulations of distributed computation [Lyn96] cannot be applied to space-time manifolds, as any manifold contains an uncountably infinite number of events, and thus the events of the computation cannot be ordered. Instead, computation may be defined in terms of fields [Bea10, BVD14, BUB13]:

Definition 1 (Continuous Computational Field)**.** *A field is a function $f : M \to \mathbb{V}$ that maps every event m in a Riemannian manifold M to some data value in $\mathbb{V}$.*

Discrete computational fields (e.g., events in the execution of a distributed algorithm) may be defined likewise, except that the domain is limited to a discrete subset of events $D \subset M$.[3]

Let $\mathbb{F}_\mathbb{V}$ be the space of all fields whose range is contained in $\mathbb{V}$. A space-time computation C is then a higher-order function that maps from each possible field defining an evaluation environment to a corresponding field of values:

Definition 2 (Space-Time Computation)**.** *A computation C is a function $C : \mathbb{F}_\mathbb{V} \to \mathbb{F}_\mathbb{V}$, where the domain of the output field is always identical to the domain of the input field.*

In other words, a computation takes a field as input, whose domain defines the scope over which the computation executes and whose values are all of the environmental state that can affect its outcome (e.g., sensor readings). At every point of space and time in the execution scope, some output value is produced.

[1] This is analogous to the physics notion of a world-line.

[2] In practice, the results we develop here often apply to slowly moving devices as well.

[3] Considering a discrete execution as a manifold subspace, rather than an abstract graph, preserves its relationship with the space in which devices are embedded.

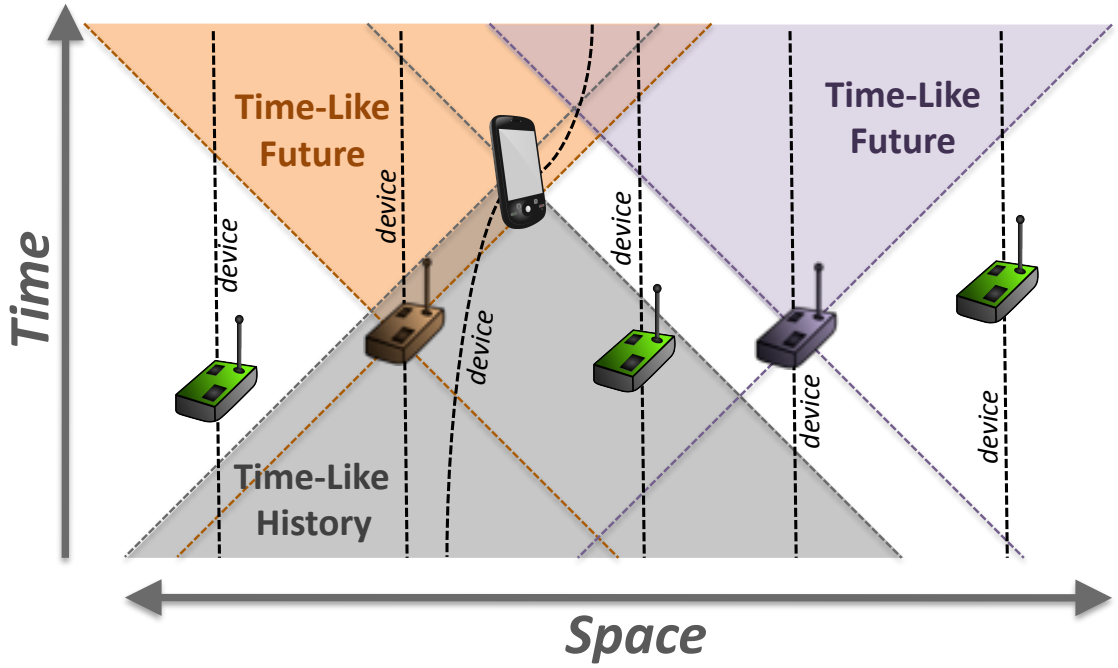

Figure 21.3: Example of space-time relations on a manifold representing a network of wireless devices. Sensor nodes are stationary, while the phone moves along a time-like trajectory (i.e., slope less than diagonal). This image resorts to a relativistic view of the world, in which the speed of light is constant and time and space are together in a continuum. In this view, an event happening far away in space becomes accessible only after some time: as an example, the light emitted from Sun takes about eight minutes to come to Earth, so what we are watching when we look at the Sun is actually its status eight minutes ago. In this image, diagonals represent light cones, namely the future and past "field of views" of events of each device.

The computations we are interested in are those physically realizable: those that are both approximable, meaning that the output field is bounded in its intricacy, and causal, meaning that information moves subject to the speed bound c. Approximability is determined by comparing values in a continuous field to the nearest points in a corresponding discrete field:

Definition 3 (ε-approximation)**.** *Let $D_\varepsilon \subset M$ be a set such that every event $m \in M$ is within distance ε of some event in D_ε. The ε-approximation of field $f : D_\varepsilon \to \mathbb{V}$ is a field mapping every point in M to the value of f at the nearest point in D_ε (choosing arbitrarily for equidistant points).*

Definition 4 (Field Approximation)**.** *A countable sequence of ε_i-approximations f_i of manifolds M_i, with $\varepsilon_i < \varepsilon_{i-1}$ is said to approximate field $f : M \to \mathbb{V}$ if both the following hold:*

$$\lim_{i\to\infty} |(M \cup M_i) - (M \cap M_i)| = 0$$

$$\lim_{i\to\infty} \int_{M\cap M_i} |f - f_i| = 0$$

In other words, a sequence of increasingly fine discrete sets approximates a continuous field if both the manifolds and the values assigned over them by the fields tend toward identical. A field produced by a computation is approximable if at least one such approximation sequence can be constructed.

The second condition for physical realizability is causality: causal computations have values that are determined only by the information available at each event. Intuitively, information is available if it could have been relayed from event to event from its origin. To compute this, we add the concept of a neighbourhood, defining how far information can move in one step:

Definition 5 (Neighbourhood)**.** *Neighbourhoods are defined by a function $\mathcal{N}$ mapping each m to a connected open set $\mathcal{N}(m) \subset M$, containing at least every event within distance r of m (for some positive r).*

Information is thus directly available to an event m in device d from the region $I(m)$ of events that have been in its neighbourhood in the present or past:

$$I(m) = \overline{T^-(m)} \cap \mathcal{N}(d - T^+(m))$$

and the set of all events from which information is available is the transitive closure of direct information availability:

$$I^\infty(m) = I(m) \cup I(I(m)) \cup \ldots$$

For continuous computations, this is precisely equal to the closure of the past (i.e., $I^\infty(m) = \overline{T^-(m)}$). For discrete computations, however, the region depends on the specifics of the discretization, because information cannot be relayed from a device except at its events. For example, Figure 21.4 shows discrete computations with the same neighbourhood but different speeds of

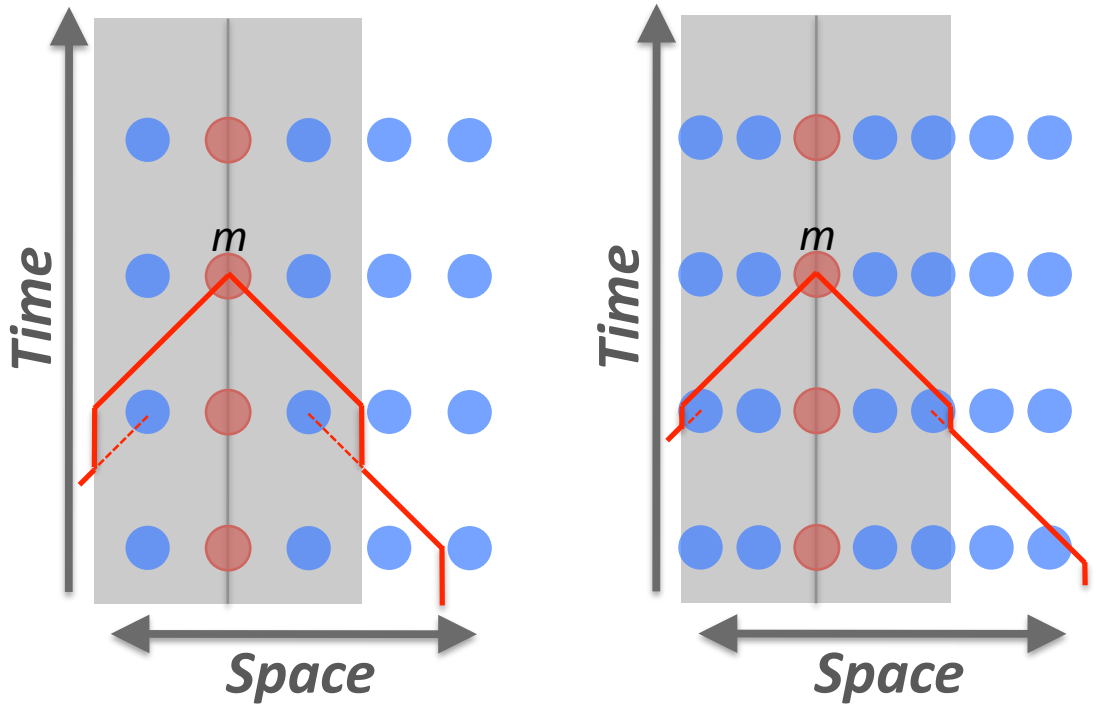

(a) Slower Discretization

(b) Faster Discretization

Figure 21.4: The effective speed of information flow for a discrete computation depends on the specifics of discretization relative to neighbourhood size. For example, in discretization (a) information moves significantly slower than in discretization (b), despite the neighbourhoods being the same, leading to a smaller information availability (red line) at event *m* of the red device.

information flow due to differences in the position of discrete events within the neighbourhood. The difference between $I^{\infty}(m)$ and $\overline{T^{-}(m)}$ is bounded by ε, however, and converges to zero with ε.

Causal computations are then those whose values depend only on events from which information is available:

Definition 6 (Causality)**.** *Computation C is causal if for every pair of evaluation environments e and e' with the same domain M, it is the case that $e(I^{\infty}(m)) = e'(I^{\infty}(m))$ implies $C(e)(m) = C(e')(m)$.*

Some examples of causal and approximable computations:

- a computation whose value is always 3;
- a computation that applies the triangle inequality to measure the spatial distance to specific events designated by the environment;
- a gossip computation that spreads knowledge that a particular environmental condition has been sensed.

An example of a non-causal computation is a computation that instantaneously changes from 0 to 1 in response to a sensor value at a single event. An example of a non-approximable computation is one that maps all rational distances from an event to 0 if the numerator is even and 1 is the numerator is odd.

21.1.3 Space-time programs

One way of specifying space-time computations is by the functional composition of a basis set of operators:

Definition 7 (Space-Time Operator)**.** *A space-time operator is a function $o : \mathbb{F}_{\mathbb{V}} \times \mathbb{F}_{\mathbb{V}}^{k} \to \mathbb{F}_{\mathbb{V}}$ taking an evaluation environment and zero or more additional fields as inputs and producing a field as output.*

This is much like the definition of a computation, except that the domains of the fields may differ.

In this work, we consider only causal operators, any functional composition of which must clearly also be causal. Some examples of causal operators:

- + produces a scalar field whose value at each point is the sum of the input values at that point.
- `true` outputs a field mapping the domain of the evaluation environment to the Boolean value *true*.
- `temperature` reads a temperature sensor by extracting its value from the evaluation environment, producing a scalar field.

A *space-time program* is then any functional composition of operator instances to form a computation, such that the domains and ranges of the output are well-defined for all possible values of the inputs [4]. Allowing compositions to be defined as new operators by usual means such as lambda calculus, programs may be recursive as well, with an evaluation structure determined dynamically by their environment. Space-time programs can be formalized as dataflow graphs[BUB13, BVD14]; here we use an equivalent syntactic formulation and consider only the restricted language developed next.

For clarity in the description of programs, we will abuse terminology; in the context of a program, a "field" is not actually the mathematical object itself, but the input or output of an operator instance, which takes on a field value when evaluated in the context of an environment.

21.1.4 Eventually Consistent Programs

We now apply the continuous model of computation to identify distributed algorithms that are resilient to changes in the number and position of devices. Our approach is as follows: consider any distributed algorithm for which we can identify an analogous space-time program. If the distributed algorithm always approximates the space-time program, no matter the particular arrangements of discrete events, then it is clearly resilient to increasing numbers of devices and to small changes in device position:

Definition 8 (Consistent Program)**.** *Let* `P` *be a space-time program, e be an evaluation environment, and* e_i *a countable sequence of* ε*-approximations that approximate field e. Program* `P` *is* consistent *if* $\texttt{P}(e_i)$ *approximates* $\texttt{P}(e_i)$ *for every e and* e_i*.*

As noted in the discussion of causality, however, information may move at different speeds in different discrete approximations. This means any program with non-trivial interaction between devices will not be consistent, but that if a program eventually converges to a steady state, then it may be consistent after that point:

Definition 9 (Eventually Consistent Program)**.** *Consider a causal program* `P` *evaluated on environment e with domain M. Program* `P` *is* eventually consistent *if, for any environment e in which there is a spatial section* S_M *such that the values of e do not change at any device in the time-like future* $T^+(S_M)$*, there is always some spatial section* S'_M *such that* `P` *is consistent on the time-like future* $T^+(S'_M)$*.*

[4] Some notes on technicalities.

- Any well-defined composition of operators is itself an operator.
- Complete programs have no inputs except the environment. Any composition of operators that requires inputs (e.g., function definitions) may, however, be transformed into an equivalent program by "currying" the inputs to instead be supplied by the environment and adjoining a special "no value" value for events in the domain of the environment but not the input.

In other words: if the inputs ever converge, then the outputs eventually converge as well, and are consistent thereafter.

As defined, eventual consistency is a conservative property, saying nothing about behaviour before convergence. For many programs, however, the constructs used to ensure eventual consistency often produce more broad resilience.

21.2 Eventually Consistent Language

In this section we develop a language that generates only eventually consistent programs. This language, GPI-calculus, derives both its syntax and semantics from Field Calculus [VDB13], but is restricted to a set of eventually consistent operators that compose to produce programs that are also guaranteed to be eventually consistent.

21.2.1 From Consistency Failures to Fragility

Consistency is a useful property because it allows us to evaluate when a computation is fragile to the circumstances of its evaluation. Even though real networks are always finite, comparison with a continuous ideal can reveal vulnerabilities in a distributed algorithm through the manner in which it fails to be eventually consistent.

Unbounded recursion, of course, can create consistency failures: infinite loops are obviously ill-defined. There are also many ways to arrange a recursion that grows in depth with density and does not converge. We can eliminate this by the simple but harsh method of prohibiting recursive function calls, which puts a finite bound on the number of operations that can be used in the evaluation of a program.

Let us turn next to a problem related to distributed computation: interactions between devices frequently lead to scalability problems due to implicit dependence on the distribution of devices. Consider, for example, the Field Calculus program:

```
(def non-convergent-nbr (v)
  (/ 1 (min-hood (nbr-range))))
```

As devices pack more closely, the distance to the closest device approaches zero, causing the inverse to rise without bound, meaning it does not converge. The program is thus not consistent, indicating its fragility to device distribution.

Likewise, state creates problems in programs that implicitly rely on the time between events:

```
(def non-convergent-rep ()
  (rep x 0 (- 1 x)))
```

This program switches values between 0 and 1 at every evaluation. As the frequency of evaluations rises, it does not settle one or the other, but oscillates faster and faster, failing to converge.

It is thus not consistent, indicating its fragility to exactly when and how frequently state updates occur.

Thus, in order to ensure resilience, we will not allow direct access to either state or neighbourhood functions, and instead embed them in higher-level constructs with guaranteed convergence.

Eliminating constructs that can directly give rise to consistency failures is not enough however. Even if a program produces approximable fields, the program may not be consistent because it may not be resilient against miniscule changes in its inputs. For example, the `distance-to` program presented in Section 3.5.2 always produces an approximable field of shortest path distances, but is not a consistent program because the value of the output may be greatly affected by individual points in the `source` field. Consider, for example, a `source` field where only one point is *true*: an ε-approximation containing that point has only finite values, while one without that point has infinity everywhere. Thus, it is possible to construct sequences that do not converge, because they alternate between including and not including the critical point.

Are such extreme inputs reasonable sources of concern, however? In fact, it is remarkably easy to produce fields with a set of critical values that has measure zero, i.e., being infinitely thin and therefore not guaranteed to be sampled in any approximation. For example, a bisecting boundary computed by comparing `distance-to` functions

```
(def bisector (point-1 point-2)
  (= (distance-to point-1) (distance-to point-2)))
```

computes a field that is *false* except at an infinitely thin boundary of *true* values. Fed to a program sensitive to such sets, such as `distance-to`, this can result in arbitrarily unpredictable behaviour from a distributed algorithm.

This is not a special case related to `distance-to`, but a deeper conflict for situated distributed algorithms, between the discrete values commonly used in algorithms (e.g., Booleans, branches, state machines) and the continuous space-time environment in which devices are embedded. In particular, any non-trivial field with a discrete range cannot be continuous, meaning that it either is itself not approximable or else contains some measure-zero boundary region that, if handled badly, can generate unpredictable behaviour (as in the bisector example).

It is not practical to eliminate every program element that could either generate boundary regions or generate unpredictable behaviour due to values at a boundary. As the bisector example shows, this cannot be done without eliminating at least one of distance measure, comparison, or branching[5], and losing any of those constructs curtails expressiveness more drastically than we are willing to accept. Instead, the approach we present accepts that problematic boundaries will exist, and instead dynamically marks them and contains their effects.

[5]Equality testing can always be implemented by any Boolean-valued comparison of scalars plus Boolean branching.

21.2.2 GPI-calculus

Having identified ways in which consistency fails, introducing fragility to distributed algorithms, we now present a new language, GPI-calculus, that restricts the syntax of Field Calculus and ensures that all programs it can express are eventually consistent. We accomplish this by dynamically marking and tracking potential problems, with the following key changes to Field Calculus:

- A new value, $\mathscr{B}$, is adjoined to the set of data values, representing a boundary between value regions, and is passed unchanged through all functions.
- Comparing real numbers produces $\mathscr{B}$, rather than a Boolean, when the numbers are precisely equal.
- State and communication are available only through a new compound operator, `GPI`, which can implement a number of distance-based operations.

With these changes, well-written programs will ensure eventual consistency by effectively excluding a miniscule (often empty) set of devices from certain computations. Poorly written programs (e.g., `(= (sqrt 2) (sqrt 2))`) may still contaminate large areas with $\mathscr{B}$ values, but will still converge—just not to a particularly useful result.

Key to this approach is the new operator `GPI`, a "gradient-following path integral," which we define as the Field Calculus function:

```
(def GPI (source initial density integrand)
 (if (<= density 0)
  ℬ // Metric ill-defined if density non-positive
  (2nd
   (rep distance-integral
    (tuple infinity initial) // Initial value
    (mux source
     (tuple 0 initial)// Source is distance zero, initial value
     (min-hood' // Minimize lexicographically over non-self nbrs
      (+ (nbr distance-integral)
       (* (nbr-range) // Scalar multiplication of tuple
        (tuple (mean density (nbr density))
         (mean integrand (nbr integrand)))))))))))
```

For this implementation, we use a slightly modified version of the usual field-calculus `min-hood` operator, designated as `min-hood'`, which returns $\mathscr{B}$ if the minimal value for the first tuple element is held by more than one device.

The `GPI` operator thus performs two tasks simultaneously. First, like `distance-to`, `GPI` computes a field of shortest-path distances to a `source` region. Here, however, distance is

```
l ::= 𝔹 | ℤ | ℝ                                   ;; Literals
b ::= m | mux | <                                ;; local operators
e ::= x | l | (b ē) | (f ē) | (sense ℤ⁺)          ;; expression
    | (if e e e) | (GPI e e e e e)               ;; special constructs
F ::= (def f(x̄) e)                               ;; function
P ::= F̄ e                                        ;; program
```

Figure 21.5: Syntax of GPI-calculus.

"stretched" proportional to a scalar field `density` (representing e.g., crowd density slowing movements, hazards increasing danger of movement). Furthermore, the `min-hood'` operator ensures that all points in the distance field with more than one shortest path are $\mathscr{B}$. Second, `GPI` computes a path integral of the scalar field `integrand` following the gradient of the distance field upward away form the source, starting at the scalar value `initial` in the source region. The function definition binds these together via a tuple and lexicographic minimization, such that the value added to the line integral at each device is taken from the neighbour on the (sole) minimal path to the source.

Importantly, the `GPI` operation subsumes a number of useful and frequently used computations. For example, `distance-to` can be computed as:

```
(def distance-to (source)
  (GPI source 0 1 1))
```

which eliminates stretch by setting it to constant 1 and obtains distance by integrating 1 along each shortest path.

The complete syntax of GPI-calculus, shown in Figure 21.5, is then defined as a restriction of Field Calculus:

- The `rep` and `nbr` constructs are only available indirectly via the `GPI` construct as defined above.
- Literals are restricted to only Booleans ($\mathbb{B}$), Integers ($\mathbb{Z}$) and real numbers ($\mathbb{R}$).
- Built-in operators are restricted to the elements `m`, `mux`, `<`, and `sense`.

The available built-in operators are:

- `m` is any strictly continuous mathematical function (e.g., addition, multiplication, logarithm, sine[6]), extended to have output $\mathscr{B}$ if any input is $\mathscr{B}$. Finally, for any m that can map

[6]This also includes construction and referencing of tuples, which can be used to implement data structures, as well as higher order functions such as map and reduce. Some simple functions are excluded, however, such as division, which is discontinuous when the denominator is zero.

integers to integers (e.g., addition), the output has integer type iff all inputs have integer type.

- `mux` is the piecewise multiplexer function: when its first input is *true*, it returns the second input; if it is *false*, it returns the third input; if it is $\mathscr{B}$, it returns $\mathscr{B}$.
- $<$ compares two numerical inputs, returning *true* if the first is less and *false* if the second is less. If they are equal, then the result depends on type: if both are integers the result is *false*; if either is a real its $\mathscr{B}$. Finally, if either input is $\mathscr{B}$, then the result is also $\mathscr{B}$.
- (`sense` k) returns the kth value in the environment state (assumed to be a tuple), where k is the positive literal given as its input.

We can interpret the same syntax with both discrete and continuous semantics; our goal will be to show that the discrete is eventually consistent with the continuous for every well-defined program (i.e., one without syntax or semantic errors).

The discrete semantics of GPI-calculus is identical to that of Field Calculus, with one restriction (function calls may not be recursive) and one trivial modification to the semantics of `if`. In Field Calculus, the first input to `if`, the test value, must be a Boolean. In GPI-calculus, it may also be $\mathscr{B}$, in which case, neither branch is evaluated, state is erased from both branches, and the result is $\mathscr{B}$. This change, however, is entirely neutral with respect to the interactions of `if` with other elements of the semantics, and so does not affect any of the properties established in [VDB13].

We may also interpret GPI-calculus with a continuous semantics [BUB13]. Given an environment domain M, literals are constant-valued fields, the built-in operators act pointwise on fields, function calls and variable references are evaluated through a substitution model, and `if` evaluates its subexpressions on environments with domains restricted to match the subspaces of its test argument. As with the discrete semantics, we prohibit recursion.

Specifying GPI-calculus as this minimal set of operations makes eventual consistency tractable to prove, yet still able to implement a much more general range of operations. For example, `mux` is sufficient to implement all logical operations, e.g.:

```
(def and (a b) (mux a b false))
(def or  (a b) (mux a true b))
(def not (x)   (mux x false true))
```

21.2.3 GPI calculus is Eventually Consistent

We can now prove that GPI calculus accomplishes the goal for which it has been designed: ensuring that any program that can be expressed using it is eventually consistent, meaning that it is resilient against the particulars of how many devices are in the network and how they are distributed in space.

Theorem 1 (Eventual Consistency of GPI-calculus)**.** *GPI-calculus programs are eventually consistent for all environments e that are continuous on* $e^{-1}(\mathbb{V} - \mathscr{B})$.

We approach this proof in three stages (see Appendix A). First, we prove that any finite composition of operators that are eventually consistent and preserve certain continuity properties is also eventually consistent and continuity preserving. We then show that each operator in GPI-calculus is individually at least eventually consistent and continuity preserving. Finally, we show that all GPI-calculus programs are equivalent to finite compositions of operators, which implies that all such programs are eventually consistent.

Thus, if a program is evaluated in a "well-behaved" environment, its results are predictable and resilient to scale and positioning of devices. Having proved this, in the next section we explore the usefulness of GPI-calculus and demonstrate its consistency properties empirically in simulation.

21.3 Example Applications

Having established a calculus for eventually consistent programs, we now aim to demonstrate its expressiveness. In this section we present a number of self-organization patterns that can be described in terms of GPI-calculus, with accompanying application scenarios in current or emerging large-scale situated networks: wireless sensor networks, ad-hoc networks created by mobile smart-devices, and urban-scale ecosystems of devices and services for crowd steering. Simulations of these scenarios demonstrate both the potential impact of our result and confirm the consistency result and its effect in providing coordination behaviors resilient to changes in network density and scale.

21.3.1 Distance-Based Patterns

Distributed distance calculation is one of the most frequently used building blocks for self-organizing systems [FMDMSM+12b]. As we have seen in the previous section, distance fields can be computed via a straightforward application of `GPI`. A second example uses the field of distance estimates as a "carrier" to broadcast by gossiping a value from the source. This can be defined from `GPI` as:

```
(def broadcast (source value)
  (GPI source value 1 0))
```

Here, `GPI` shifts the initial `value` outward by integrating 0 along the path up the distance field, so that no matter how far away, the value is always the same. When the source region has a single value v, this produces a field that stabilizes to v at every device, intuitively a "broadcast." In the more general case where different devices in the source have different values, the resulting field maps each device to the value of `value` in the nearest source device.

Functions `distance-to` and `broadcast` can be functionally composed to create a higher-level `channel` distributed structure useful for tasks such as corridor routing. This version of channel takes two source fields `a` and `b`, and a (typically constant) numerical field `w`, and yields a boolean field holding `true` only in those devices whose distance to `a` and `b` is no more than `w` greater than the shortest path:

```
(def channel (a b w)
  (< (+ (distance-to a) (distance-to b))
     (+ w (broadcast a (distance-to b)))))
```

The `broadcast` expression sends to every device the distance from `b` perceived by devices in `a`. Composition by the built-in functions < and + does the remainder of the job. This example emphasizes how function composition can generate a plethora of spatial structures [MPV12a].

Another key feature of GPI-calculus is the ability to modulate the behavior of algorithms by restricting their domain with construct `if`. For example, this can be used to create channels that circumvent a given area considered as an obstacle:

```
(def channel-with-obstacle (a b w obstacle)
  (if obstacle false (channel a b w)))
```

Where `obstacle` is `true`, there is no channel and `false` is returned; in the remainder of the network, we compute a channel in the usual way and manifold geometry routes it around the "missing" space.

As an application scenario, consider a wireless sensor network that is heterogeneous distributed in two dimensions and in which some devices much exchange a large amount of information, e.g., relaying video of an ongoing event to a mobile monitoring station. The goal is to identify a set of devices that can perform the relay, such that:

- the information spreads to only a small fraction of devices (e.g., to save battery energy);
- there is some replication along the transmission path, to reduce the probability of data loss;
- if some devices do not want to collaborate to the transfer (e.g. because their battery level is low, or because they have shown faults), they should be circumvented.

The `channel-with-obstacle` pattern provides one possible implementation, marking only a set of devices near the shortest path excluding non-participating devices. A `broadcast` restricted to this channel with `if` can then effectively relay the data with enough replication to prevent data loss, but with much less resource consumption than an unrestricted broadcast. Figure 21.6 shows an example of this scenario simulated in ALCHEMIST [PMV13]. Note how the consistency of the pattern does not change significantly for different densities of devices: a change in density only affects the precision of the channel shape.

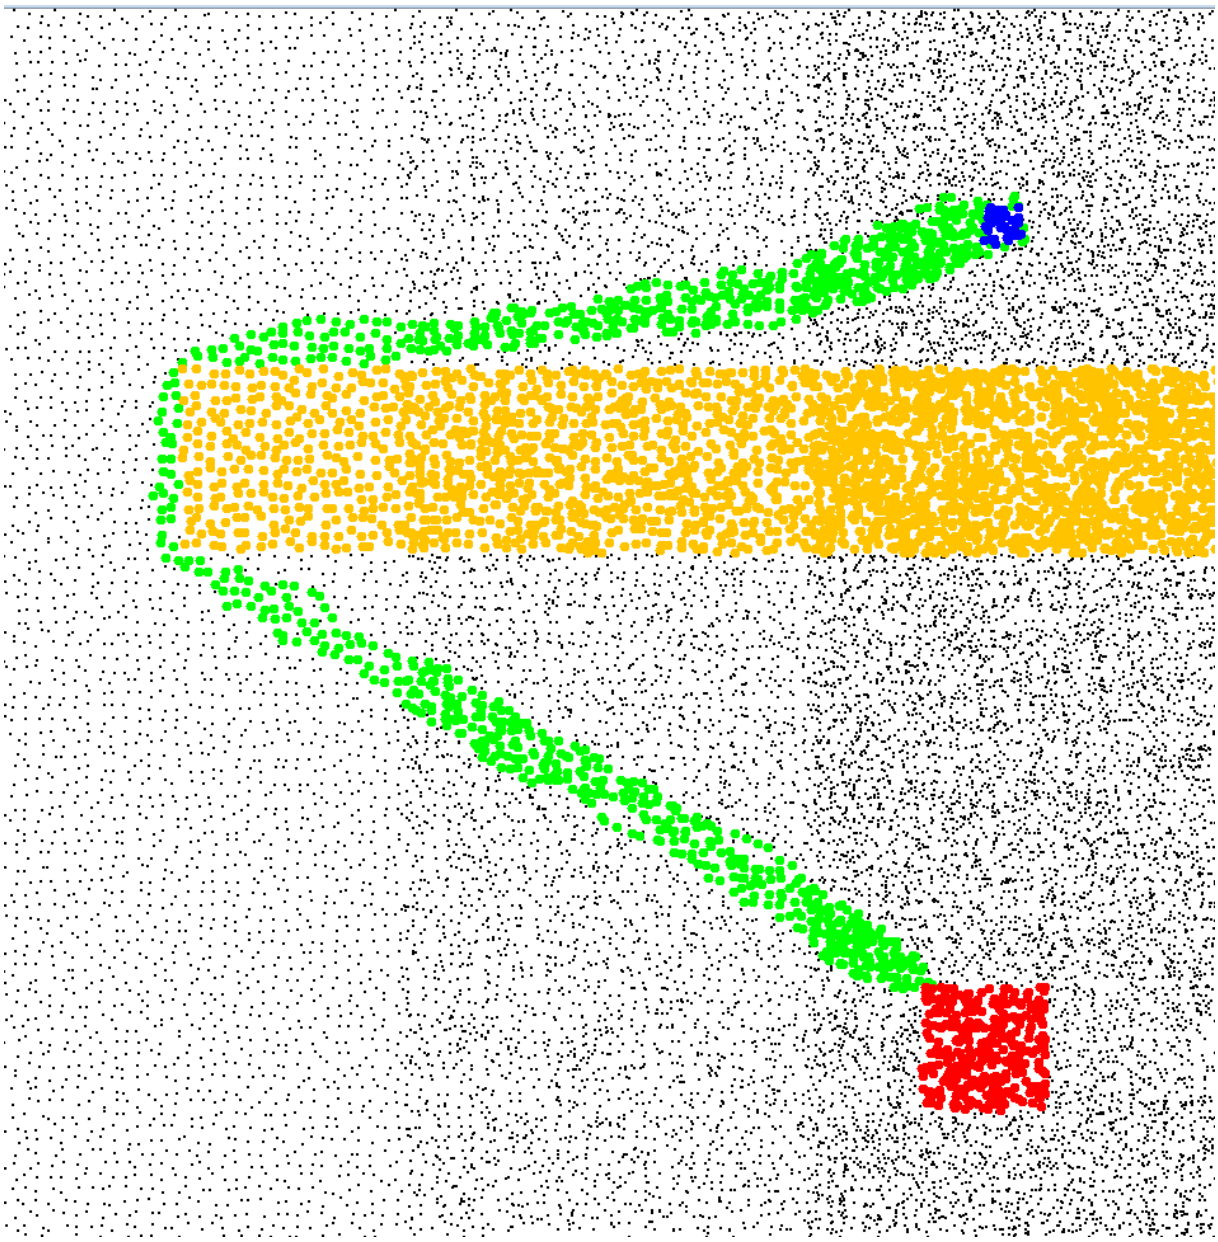

Figure 21.6: Channel pattern deployed on a wireless sensor nework with 20,000 devices deployed at heterogeneous density. Blue devices send data to red along the channel (green), avoiding the obstacle (yellow).

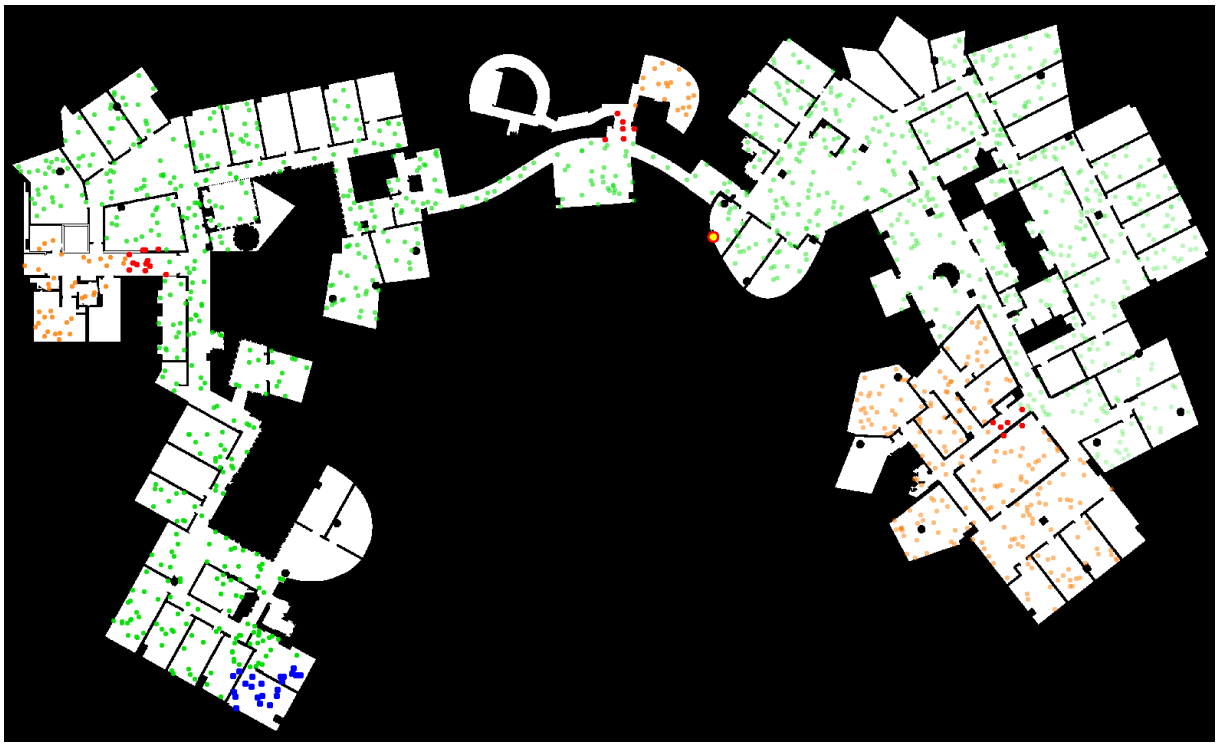

Figure 21.7: Obstacle forecast pattern deployed to alert people moving through a building of obstacles (red) on their path to a destination (blue). Distances are shown in green, and alerted areas where people must be advised of the obstacle are orange.

21.3.2 Path Forecasting

In the uses of the `GPI` construct so far the `integrand` argument has been constant at 1 for measuring distance and 0 for broadcast. Most generally, however, this argument to `GPI` can be used to record the context along a path, by integrating the values encountered along the way.

An example of how this can be used is to forecast obstacles that may be encountered along a path, and provide warnings thereof. For example, if we set `integrand` to an indicator field that is 1 where there is an obstacle and 0 where there is not, then `GPI` outputs a field that holds a positive value only on those devices that are reached by crossing an obstacle:

```
(def obstacle-forecast (source obstacle)
  (GPI source 0 1 obstacle)
```

As an application scenario, consider a building occupied by a number of people, who need to navigate through it with various goals, e.g., finding a desired location, rendezvous with another person, exiting while avoiding hazards. Some areas of the building are currently problematic and should be avoided, e.g., due to crowding or temporary barriers or dangerous hazards. The

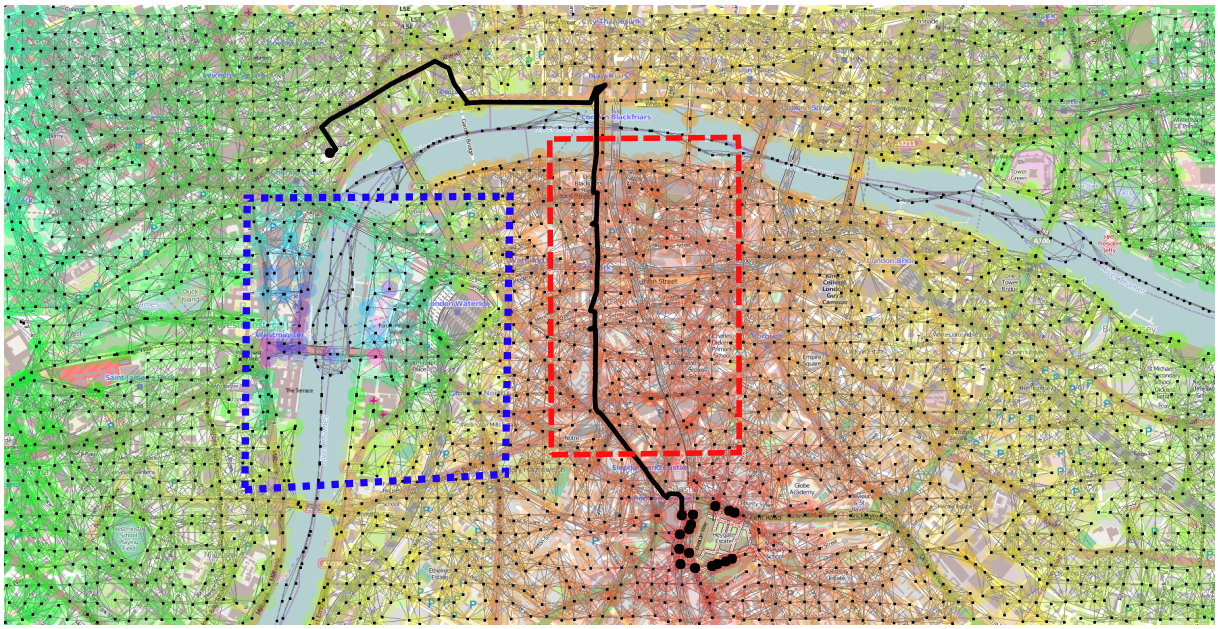

Figure 21.8: Context-sensitive distance computation deployed for navigation on a street map of London. Warmer coloured devices have a closer effective distance to the destination (black dots at bottom centre). Dashed rectangular areas represent unfavourable (blue) and favourable (red) areas for travel. An example path is shown (black line), originating near Charing Cross (black dot in upper left).

goal is then to alert people who would normally transit across those areas while moving to their destination.

The `obstacle-forecast` pattern is a possible solution for this problem. Applying it with regards to a given destination (used, ironically, as the `source` argument), the alerted areas are those where `obstacle-forecast` returns a value greater than zero.

Figure 21.7 shows an example of this scenario simulated in ALCHEMIST for the MIT Stata Center floorplan from Figure 21.2. Note the complex manifold structure of the building modifies the shape of the sector pattern: the presence of walls considerably changes the distances compared to an open area. Despire the complexity, however, the manifold geometry ensures an eventually consistent pattern. Note also that, having alerted people to obstacles, the system could go on to provide alternate routes by running `distance-to` on a space omitting the obstacles.

21.3.3 Context-Sensitive Distance

The effective shortest path between a node and the source of a `GPI` is not necessarily the physically shortest path induced by the standard metric. Rather, the notion of effective distance may be influenced by other properties of the environment, either negatively (e.g., obstacles, congestion, pollution, tolls) or positively (e.g., safety, dedicated lanes, beauty). The `density`

argument of GPI allows such factors to be taken into consideration as a multiplicative "stretching" of the base physical distance metric. Assuming there are penalising areas (cons) and favourable areas (pros), both expressed as scalar fields with values between 0 (least significant) and 1 (most significant), then one way to define context sensitive distance for use in navigation is:

```
(def contextual-distance (source pros cons)
  (GPI source 0 (+ 1.1 (- cons pros))))
```

As an application scenario, consider guiding pedestrian or vehicle traffic in a complex urban environment. Devices are deployed along and around the streets. Some have the ability to sense environmental parameters such as crowding, current and expected traffic, and pollution; other parameters, such as presence of events and attractions or comments on beauty of an area, are drawn from databases of local information, either remote or distributed and localized in the network.

From these, the devices can compute the pro and con fields for contextual distance for people navigating through the city, building a spatial structure that alters the perceived distances toward a location by taking desirability into account, so that a longer path crossing desirable areas will be favored over a somewhat shorter path crossing undesirable areas.

Figure 21.8, we shows an example of this scenario for the center of London, simulated in ALCHEMIST. Two areas are marked as favorable and unfavorable. Distances are shown with respect to a destination on the South side of the Thames, showing significant asymmetries caused both by the shape of the city and the marked areas acting as an attractor and repulsor of pedestrians, respectively. In this scenario both a complex environment and a non-physical distance metric are used, and still provide eventually consistent results.

21.3.4 Confirmation of Results in Simulation

As a confirmation of our results for GPI-calculus and its application to the described scenarios, we have simulated each scenario on a wide range of densities of devices using ALCHEMIST. Based on previous analysis, we expect that for each scenario, the values yielded by devices will converge to a stable set of values that no longer change. Furthermore, as the number of devices increases, the values converged to will themselves converge as the network of devices more closely approximates a continuous space.

Conditions for the scenario simulations are as follows.

- Each scenario is run with five logarithmically scaled densities, ten runs per condition: the outdoor scenarios with 100 to 10,000 devices, the indoor scenario with 100 to 1000 devices.

- Devices are distributed randomly through the portion of the space not blocked by environmental obstacles.

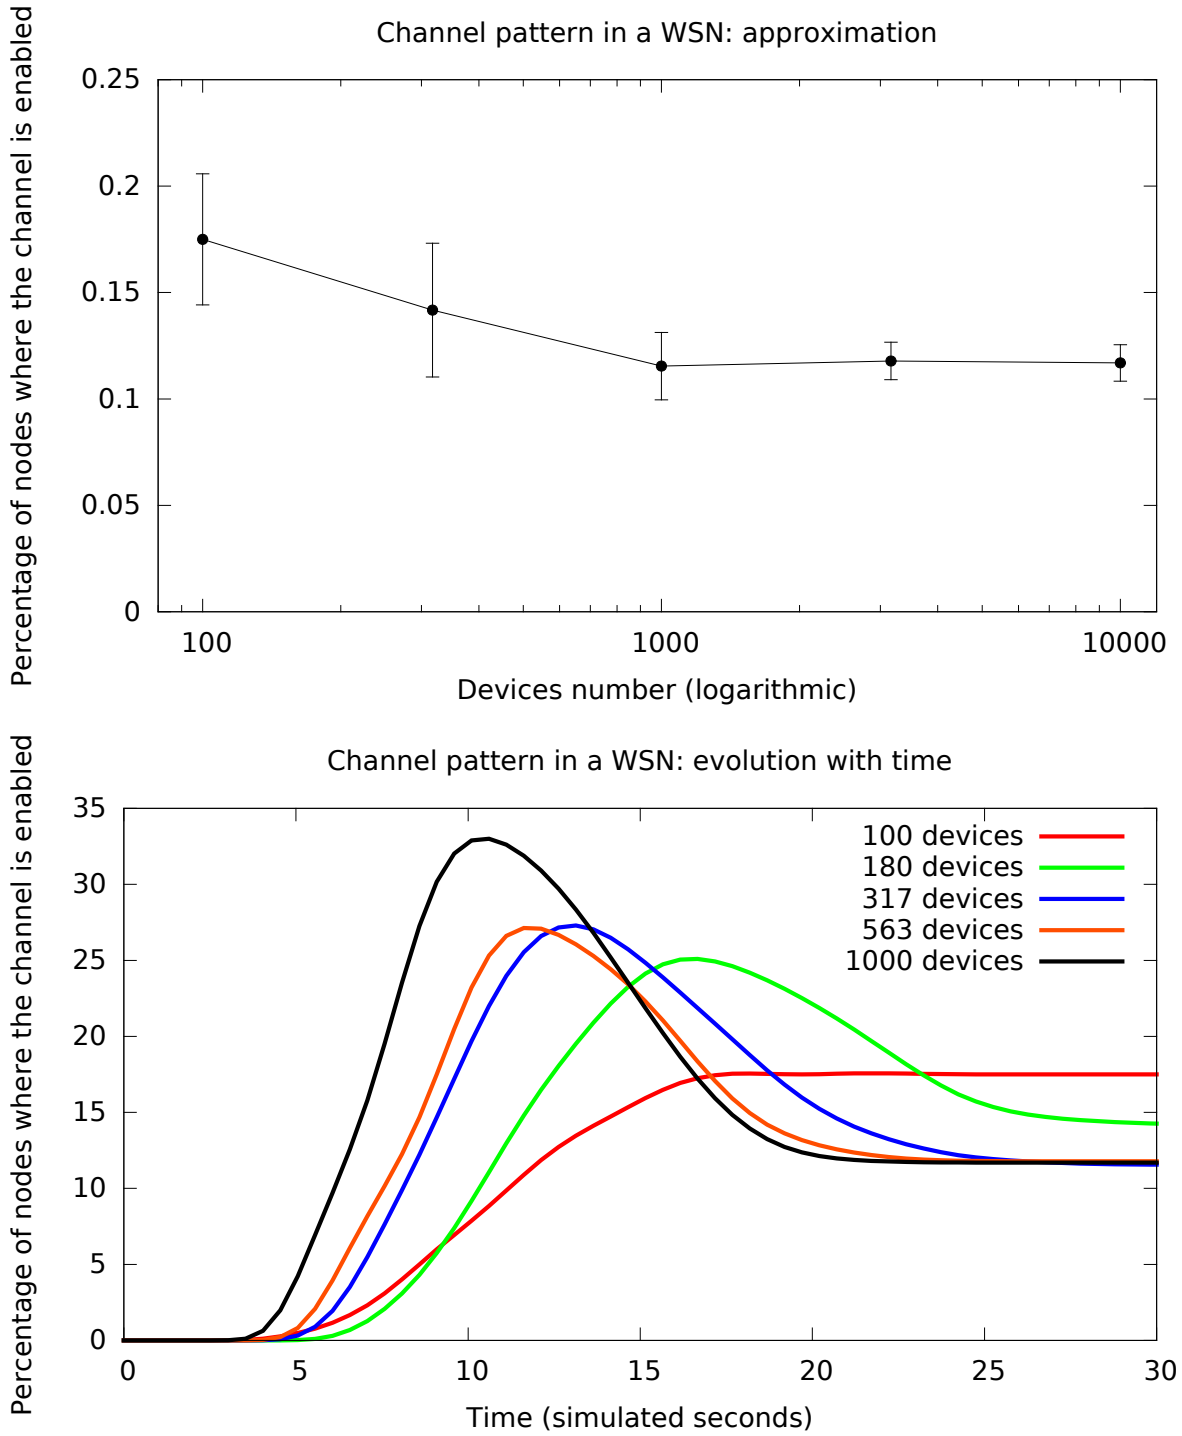

Figure 21.9: Wireless Sensor Network

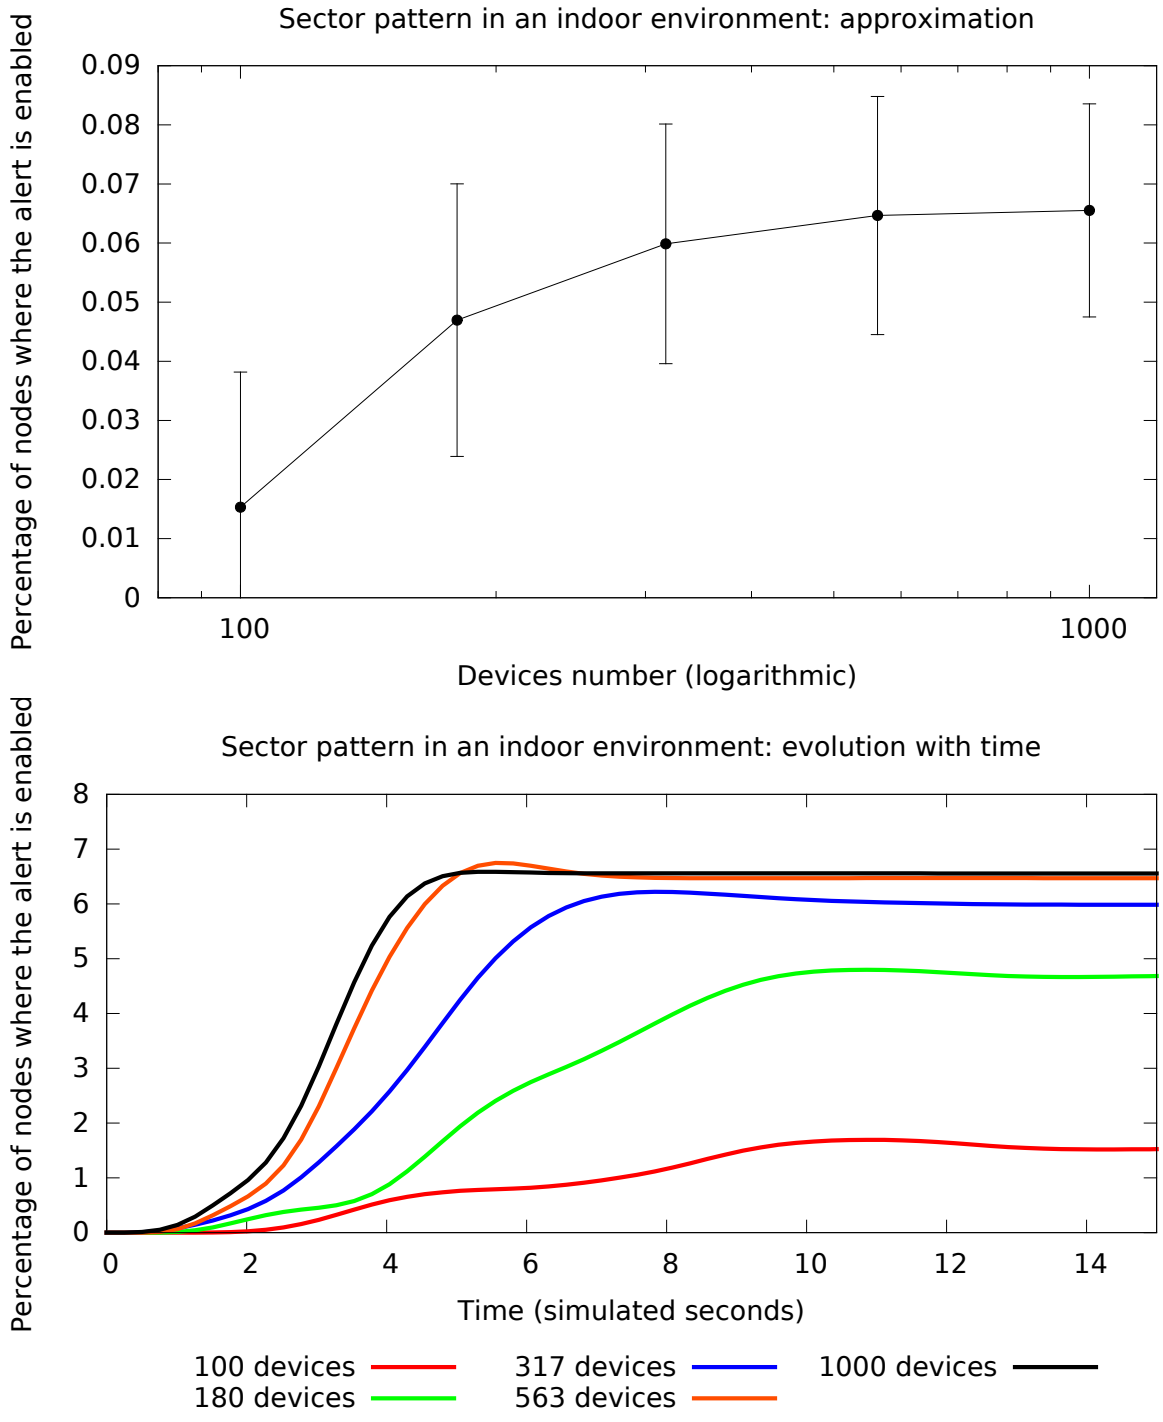

Figure 21.10: Building Alert

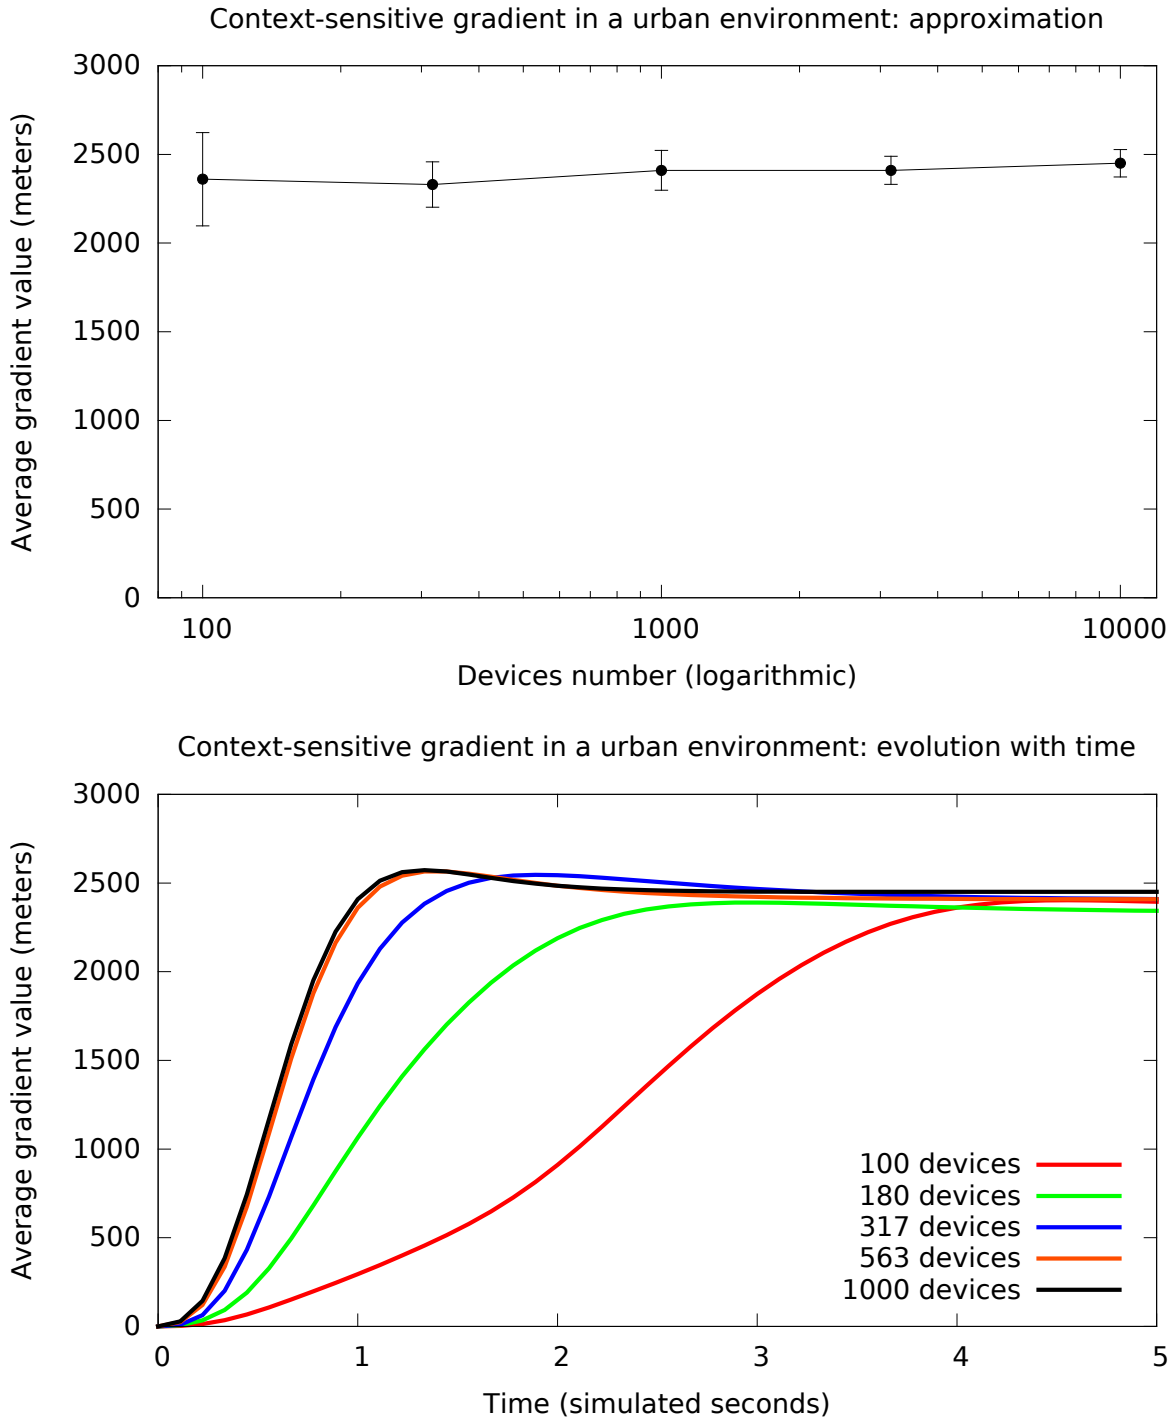

Figure 21.11: Urban Traffic Steering

- Communication range is set to 15% of width for the lowest density and reduced proportional to the square root of density, to ensure a consistent expected number of neighbours. Devices are connected with a unit disc rule, with permissive line-of-sight blocking by obstacles. In particular, two devices are connected if they are within distance r_1 and there is some position within a small distance $r_2 \propto r_1$ of each device in which no object blocks line-of-sight.

- Devices execute independently following a Poisson distribution, with frequency 1 Hertz for the lowest density; to compensate for shorter hops, frequency rises inversely proportional to communication range.

Results are summarized in Figure 21.9, Figure 21.10, and Figure 21.11, with sample snapshots shown in Figure 21.6, Figure 21.7, and Figure 21.8. For each set of simulations, we measure the mean and standard deviation of a key application property with respect to both density and time. The properties measured are:

- wireless sensor network: fraction of devices in channel;

- building alert: percentage of devices receiving the alert;

- urban traffic steering: average compensated distance value.

As predicted from our analytical results, these values converge with respect to both time and number of devices, confirming our predictions.

22 Protelis: practical aggregate programming

We have designed the Protelis language as an implementation of the Field Calculus [VDB13] closely related to Proto [BB06]. On the one hand, it incorporates the main spatial computing features of the Field Calculus, hence enjoying its universality, consistency, and self-stabilization properties [BVD14, VD14]. On the other hand, it turns the Field Calculus into a modern specification language, improving over Proto by providing

- access to a richer API through Java integration;
- support for code mobility through first-order functions;
- a novel syntax inspired by the more widely adopted C-family languages.

Protelis is freely available and open source, and can be downloaded as part of the ALCHEMIST distribution.

22.1 Syntax

We present the Protelis language in terms of its abstract syntax, provided in Figure 22.1 as a means to guide the discussion of the language's features. This syntax uses similar conventions to well-known core languages like Featherweight Java [IPW01]. We let meta-variable `f` range over names of user-defined functions, `x` over names of variables and function arguments, `l` over literal values (Booleans, numbers, strings), `b` over names of built-in functions and operators (including the "hood" functions described in Section 22.3), `m` over Java method names, and `#a` over aliases of static Java methods. All such meta-variables are used as non-terminal symbols in Figure 22.1. Overbar notation $\overline{y}$ generally means a comma-separated list $y_1, \ldots, y_n$ of elements of kind y, with the two exceptions that in $\overline{\texttt{F}}$ we use no comma separator, and in $\overline{\texttt{s}}$`;` semi-colon is used as separator instead.

```
P ::= Ī F̄ s̄;                                   ;; Program
I ::= import m | import m.*                     ;; Java import
F ::= def f(x̄) {s̄;}                             ;; Function definition
s ::= e | let x = e | x = e                     ;; Statement
w ::= x | l | [w̄] | f | (x̄)->{s̄;}               ;; Variable/Value
e ::= w                                         ;; Expression
    | b(ē) | f(ē) | e.apply(ē)                  ;; Fun/Op Calls
    | e.m(ē) | #a(ē)                            ;; Method Calls
    | rep(x<-w){s̄;}                             ;; Persistent state
    | if(e){s̄;}else{s̄′;}                        ;; Exclusive branch
    | mux(e){s̄;}else{s̄′;}                       ;; Inclusive branch
    | nbr{s̄;}                                   ;; Neighbourhood values
```

Figure 22.1: Protelis abstract syntax, colored to emphasize definition and application (red), functions (blue), variables (green), and special Field Calculus operators (purple).

22.2 Ordinary Language Features

One of the distinctive elements of Protelis when compared to other aggregate programming languages (particularly Proto), is the adoption of a familiar C- or Java-like syntax, which can significantly reduce barriers to adoption. Despite this syntactic similarity, Protelis is a purely functional language: a program is made of a sequence of function definitions ($\mathtt{F}_1 \ldots \mathtt{F}_n$), modularly specifying reusable parts of system behavior, followed by a main block of statements. Following the style of C-family languages, a function's body is a sequence of statements surrounded by curly brackets. As in the Scala programming language[1], however, statements can also be just expressions, and a statement sequence evaluates to the result of the last statement. Each statement is an expression to be evaluated (`e`), possibly in the context of the creation of a new variable (`let x = e`) or a re-assignment (`x = e`)[2]. As an example, consider the following function taking four fields as parameters, after the "channel" pattern [But02]:

```
def channel(distA, distB, distAB, width) {
    let  d  =  distA + distB;
    d  =  d - distAB;
    d < width
}
```

[1] http://www.scala-lang.org.

[2] Technically, "re-assignment" is actually the creation of a new variable that shadows the old.

This function assumes that its input `distA` maps each device to its distance to a region *A*, `distB` maps each device to its distance to a region *B*, `distAB` is a constant field holding at each device the minimum distance between regions *A* and *B*, and `width` is a constant field holding at each device the same positive number. The function then computes a Boolean field, mapping each device to `true` only if it belongs to a "channel" area around the shortest path connecting regions *A* and *B* and approximately `width` units wide. All devices elsewhere map to `false`.

Atomic expressions `w` can be literal values (`l`), variables (`x`), tuples ($[\overline{w}]$), function names (`f`) or lambdas ($(\overline{x})$ `->{`$\overline{s}$`;}`). Structured expressions include three kinds of "calls": *(i)* `b`$(\overline{e})$ is application to arguments $\overline{e}$ of a built-in operation `b`, which could be any (infix- or prefix-style) mathematical, logical or purely algorithmic function [3] ; *(ii)* `f`$(\overline{e})$ is application of a user-defined function; and *(iii)* `e.apply`$(\overline{e})$ is application of arguments to an expression `e` evaluating to a lambda or function name. The following shows examples of such calls:

```
def square(x) {
   x * x;
}
let f = square;
let g = (x) -> {
   square(x) + 1
};
f.apply(g.apply(2))        // gives 25 on all devices
```

In addition, arbitrary Java method calls can be imported and used by Protelis: *(i)* `e.m`$(\overline{e})$ is method call on object `e` and *(ii)* `#a`$(\overline{e})$ is invocation of a static method, via an alias `#a` (always starting with '#') defined by an `import` clause. The alias is created automatically as the bare method name for single imports or imports of all methods in a class with `*`. Protelis can thus interact with Java reflection to support dynamic invocation of arbitrary Java code, as shown in the following example:

```
import java.lang.Class.forName
let c = #forName("String");
let m = c.getMethod("length");
m.invoke("Lorem ipsum dolor sit amet")// gives 26 on all devices
```

22.3 Special Field Calculus Operators

The remaining constructs of Protelis are the special operations specific to Field Calculus, dealing with the movement of information across space and time:

[3] For simplicity of presentation, we omit the syntax for infix operations and order of operations, which is closely patterned after Java.

- Construct `rep(x<-w){`$\overline{\texttt{s}}$`;}` defines a locally-visible variable `x` initialized with `w` and updated at each computation round with the result of executing body `{`$\overline{\texttt{s}}$`;}`: it provides a means to define a field evolving over time according to the update policy specified by `{`$\overline{\texttt{s}}$`;}`.

- Construct `nbr{`$\overline{\texttt{s}}$`;}` executed in a device gathers a map (actually, a field) from all neighbours (including the device itself) to their latest value from computing $\overline{\texttt{s}}$. A special set of built-in "hood" functions can then be used to summarize such maps back to ordinary expressions. For example, `minHood` finds the minimum value in the map.

- The branching constructs `mux(e){`$\overline{\texttt{s}}$`;}else{`$\overline{\texttt{s}}'$`;}` and `if(e){`$\overline{\texttt{s}}$`;}else{`$\overline{\texttt{s}}'$`;}` perform two critically different forms of branching. The `mux` construct is an inclusive "multiplexing" branch: the two fields obtained by computing $\overline{\texttt{s}}$ and $\overline{\texttt{s}}'$ are superimposed, using the former where `e` evaluates to `true`, and the second where `e` evaluates to `false`. Complementarily, `if` performs an exclusive branch: it partitions the network into two regions: where `e` evaluates to `true` $\overline{\texttt{s}}$ is computed, and elsewhere $\overline{\texttt{s}}'$ is computed instead.

The following code shows some example uses of these constructs:

```
def count() {
   rep(x<-0){ x + 1}
}

def maxh(field) {
   maxHood(nbr{field})
}

def distanceTo(source) {
   rep(d <- Infinity) {
      mux (source) {
         0
      } else {
         minHood(nbr{d} + nbrRange)
      }
  }
}

def distanceToWithObstacle(source,obstacle) {
   if (obstacle) {
      Infinity
   } else {
      distanceTo(source)
   }
}
```

Function `count` yields an evolving field, counting how many computation rounds have been executed in each device. Function `maxh` yields a field mapping each device the maximum value of `field` across its neighbourhood—note a `nbr` construct should always be eventually wrapped inside a "hood" function. Function `distanceTo` nests `nbr` inside `rep` to create a chain of interactions across many hops in the network, computing minimum distance from any device to the nearest "source device" (i.e., where `source` holds `true`). It does so by a field `d` initially `Infinity` everywhere, and evolving as follows: `d` is set to 0 on sources by `mux`, and elsewhere takes the minimum across neighbours of the values obtained by adding to `d` the estimated distance to the current device—a triangle inequality relaxation computing a distance field also often termed *gradient* [LK87, BBVT08, VCMZ11]. Finally, function `distanceToWithObstacle` shows exclusive branch at work; `distanceTo(source)` is computed in the sub-region where there is no obstacle, which causes the computation of distances to implicitly circumvent such obstacles.

22.4 Architecture

In Protelis, we designed an architecture, subsumed in Figure 22.2(a), following the same general pattern as was used for the Proto Virtual Machine [BB07]. A similar architecture can also be found in The ONE-SAPERE, a simulator developed to let the real SAPERE middleware run in a controlled environment [FAS+14]. First, a parser translates a text Protelis program into a valid representation of Field Calculus semantics. This is then executed by a Protelis interpreter at regular intervals, communicating with other devices and drawing contextual information from environment variables implemented as a tuple store of $(token, value)$ pairs. This abstraction is instantiated for use on particular devices or simulations by setting when executions occur, how communication is implemented and the contents of the environment.

We have chosen to implement this architecture in Java. One key reason for this choice is that Java is highly portable across systems and devices. Another key reason (discussed further in the next section) is that Java's reflection mechanisms make it easy to import a large variety of useful libraries and APIs for use in Protelis. Finally, the pragmatics of execution on embedded devices have also changed significantly since the publication of [BB07]: a much wider variety of low cost embedded devices are now capable of supporting Java, while at the same time improvements in Java implementations have made it much more competitive in speed and resource cost with low-level languages like C [BSB+03, ORAI11].

In particular, we have chosen to implement Protelis and its architecture via the Xtext language generator [EB10] and within the ALCHEMIST framework [PMV13]. Usefully, Xtext also features support for generating a language-specific Eclipse plug-in, which provides developer assistance through code highlighting, completion suggestions, and compile-time error detection.

For an initial validation, we have exercised this architecture by construction of two instantiations: one in the ALCHEMIST framework for simulation of large-scale spatially-embedded systems; the other as a daemon for coordinating management of networked services. Figure 22.2(b) shows the ALCHEMIST instantiation: simulations are configured using a simple scripting language, which specifies a Protelis program as well as the collection of devices that will execute it, communication between those devices, and other aspects of the environment to be simulated. The ALCHEMIST event-driven simulation engine then handles execution scheduling, message delivery, and updates to the environment tuple store. Figure 22.2(c) shows the network service management instantiation. Here, each Protelis device lives on a separate server in an enterprise network, and is tethered to the networked service it is intended to manage by a service manager daemon. This daemon monitors the service, injecting information about its status and known dependencies into the environment and maintaining a neighbourhood by opening parallel communication links to the corresponding daemons on any other servers that the monitored service communicates with. Examples using each of these implementations are shown in Section 22.5 and Section 22.6.

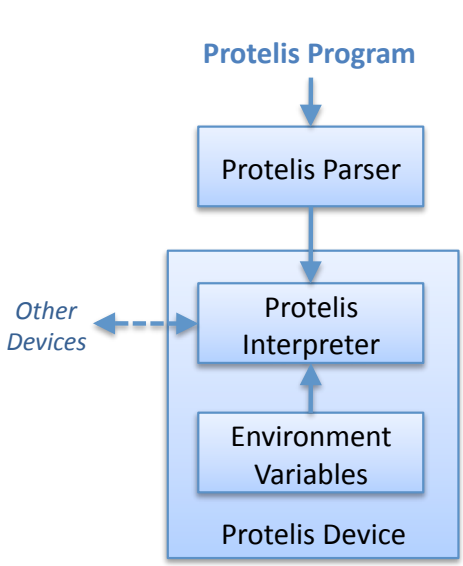

(a) Abstract Architecture

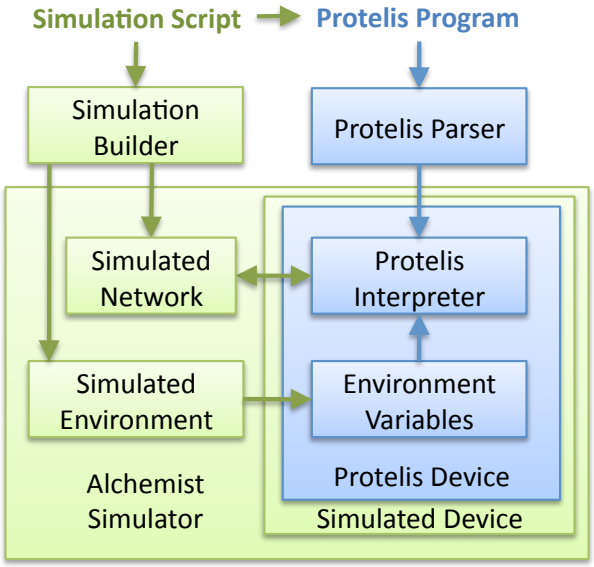

(b) ALCHEMIST Simulation

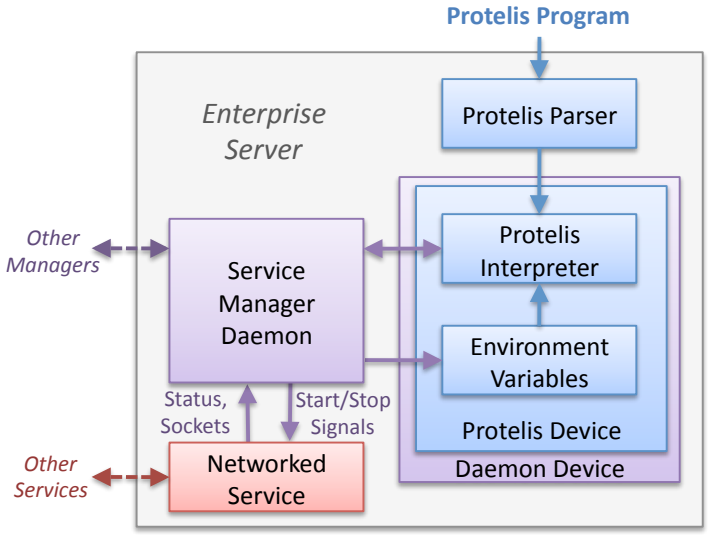

(c) Network Service Management

Figure 22.2: In the abstract Protelis architecture (a), an interpreter executes a pre-parsed Protelis program at regular intervals, communicating with other devices and drawing contextual information from a store of environment variables. This is instantiated by setting when executions occur, how communication is implemented and the contents of the environment. Two such instantiations are presented in this thesis: as a simulation in the ALCHEMIST framework (b) and as a daemon for coordinating management of networked services (c).

22.5 Application example: Rendezvous at a Mass Event

A common problem in large public events is to rendezvous with other companions attending the same large public event. At mass events, access to external cloud-based services may be difficult or impossible, and pre-arranged rendezvous points may be inaccessible or inconveniently distant. Simple peer-to-peer geometric calculations across the network, however, can readily compute a route that will allow two people to rendezvous:

```
// Follow the gradient of a potential field down from a source
def descend(source,potential) {
    rep(path <- source) {
        let nextStep = minHood(nbr([potential, self.getId()]));
        if (nextStep.size() > 1) {
            let candidates = nbr([nextStep.get(1), path]);
            source || anyHood([self.getId(), true] == candidates)
        } else {
            source
        }
    }
}
def rendezvous(person1, person2) {
    descend (person1 == owner, distanceTo(person2 == owner))
}
// Example of using rendezvous
rendezvous("Alice", "Bob");
```

Figure 22.3 shows an example of running this rendezvous process in a simulated city center. We chose London as a simulation environment, using ALCHEMIST's capability for importing OpenStreetMap data. We displaced 1000 devices randomly across the city streets (represented by pale blue dots), with a communication range of 475 meters (this range chosen to ensure no network segmentation). We then picked two devices whose owners want to meet: one device on Lambeth Bridge (lower left of the image) and one device on Tower Bridge (upper right), each marked with a yellow square. To mark the devices for Protelis, we injected their environments with a property `owner`, assigning the strings `"Alice"` and `"Bob"` as values for the first and the second device respectively.

Implementing this application requires only 21 lines of code: the listing above and the function `distanceTo` that can be found in Section 22.3. This implementation measures distance to one of the participants, creating a potential field, then, starting from the other one, builds an optimal path descending the distance potential field to return to the first participant at distance zero. The first half of the algorithm has already been described, and relies on `distanceTo`, while the second half is implemented by the function `descend`. This function,

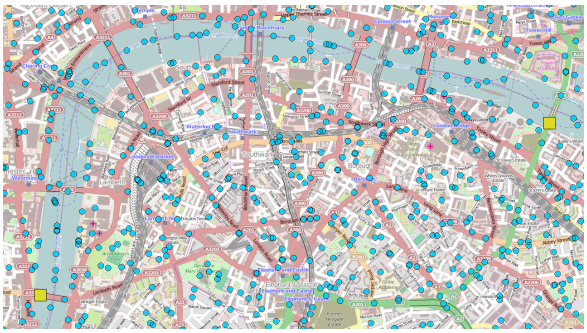

(a) Initial configuration

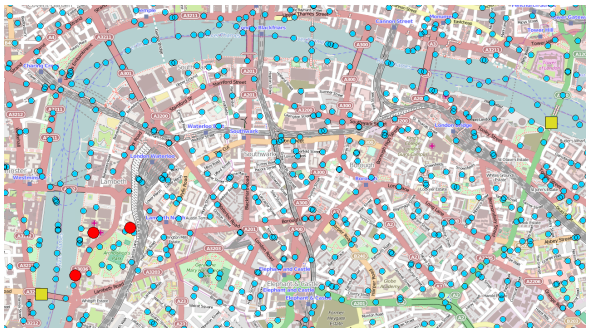

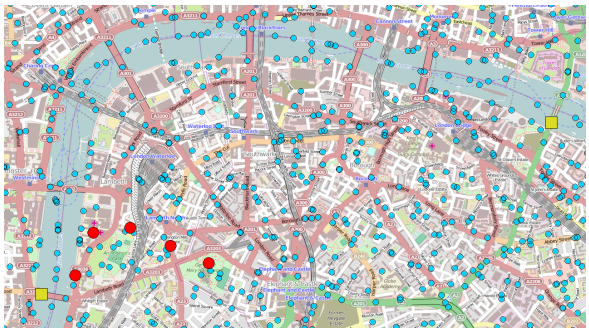

(b) Path begins to form

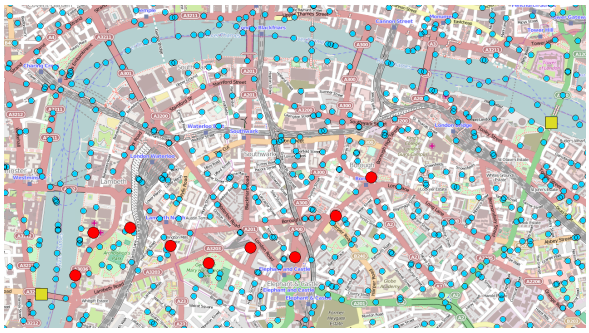

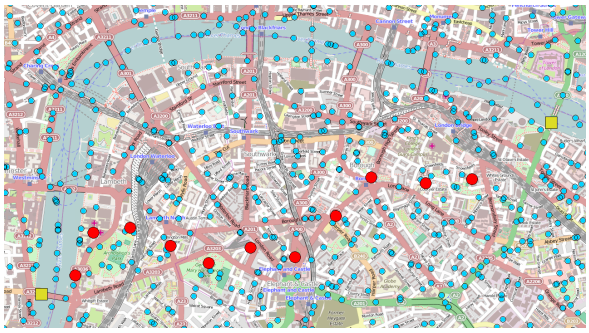

(c) Path continues to extend

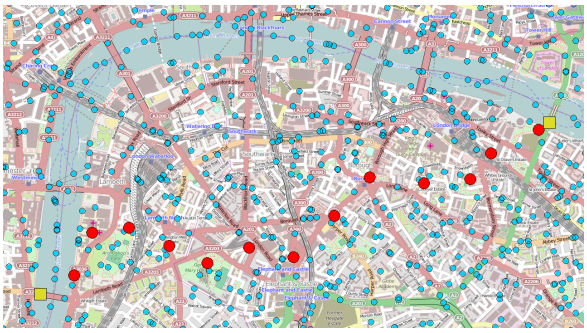

(d) Path computation complete

Figure 22.3: Example of computing a rendezvous route for two people in a crowded urban environment.

given a device and a potential field, builds a path of devices connecting the former with the source of the latter. The strategy is to mark the device we want to connect to the potential field's source as part of the path, and then, in every device, compute which of the neighbours is closest to the destination. Given this information, a device is in the path if one of the neighbours is in the path already and has marked this device as the closest of its neighbours towards the destination. Note how the whole algorithm can be elegantly compressed into just a few lines of code, and how there is no need to explicitly declare any communication protocol for exchanging the required information, thanks to the repeated use of the `nbr` operator.

As Figure 22.3 shows, once the simulation starts, a chain of devices is rapidly identified (red dots), marking a sequence of way-points for both device owners to walk in order to meet in the middle. Note also that, due to the ongoing nature of the computation, if one of the device owners moves in a different direction instead, the path will automatically adjust so that it continues to recommend the best path for rendezvous.

22.6 Application example: Network Service Management

One of the common problems in managing complex enterprise services is that there are often many dependencies between different servers and services. Frequently, some of these services are legacy or poorly coded, such that they do not respond gracefully to the failure of their dependencies. These services may continue to attempt to operate for some time, creating inconsistent state, or may be unable to resume service correctly after the server they depend on is brought back on line.

Thus, responding to a service failure often requires a coordinated shutdown and restart of services in an order dictated by service dependencies. This type of service management can be automated by attaching a daemon that watches the state of each service, then communicates with the daemons of other services to coordinate shutdown and restart in accordance with their dependencies.

Figure 22.4(a) shows an example scenario of an enterprise network for a small manufacturing and supply company, with dependencies between two key databases and the internal and external servers running web applications. This scenario was implemented on a network of EmuLab [WLS+02] servers. The services were emulated as simple query-response networking programs in Java that entered a "hung" state either upon being externally triggered to crash or after their queries began to consistently fail.

Each service was wrapped with an embedded Protelis execution engine, which was interfaced with the services by a small piece of monitoring glue code that inserted environment variables containing an identifier for the `serviceID` running on that server, a tuple of identifiers for `dependencies`, and the current `managedServiceStatus` of `stop`, `starting`, `run`, `stopping`, or `hung`. The glue code also provides `stopService` and `startService` methods to send signals to the service, tracks interactions between the services in order to maintain the set of neighbours for Protelis, and allows an external monitoring application to

attach and receive status reports.

Dependency-directed coordination of service starting and stopping was then implemented as follows:

```
import it.unibo.alchemist.language.protelis.datatype.Tuple.*
import com.bbn.a3.distributedrestart.DaemonNode.*

// Compare required and available services
let nbr_set = unionHood(nbr([serviceID]));
let nbr_missing = dependencies.subtract(nbr_set);
let nbr_required = #contains(dependencies,nbr(serviceID));
let nbr_down = nbr(managedServiceStatus=="hung" ||
                   managedServiceStatus=="stop");

// Is service currently safe to run?
let problem = anyHood(nbr_down && nbr_required) ||
              !nbr_missing.isEmpty();

// Take managed service up and down accordingly
if (managedServiceStatus=="run" && problem) {
  #stopProcess(managedService);
} else {
  if (managedServiceStatus=="stop" && !problem) {
    #startProcess(managedService);
  } else {
    managedServiceStatus
  }
}
```

In this program, each device shares information about its service ID and status with its neighbours, enabling them to track which dependencies are currently down or missing. When there is a problem with dependencies, the device invokes `stopProcess` to shut its service down, when dependencies are good, it brings it up again with `startProcess`, and when it is hung it waits for a human to sort out the problem.

Figure 22.4(b) shows a typical screenshot of the network of services in operation on an EmuLab network of Ubuntu machines, one service per machine, as visualized by the monitoring application. In this screenshot, the supplies database has crashed, causing many of the other services to gracefully shut themselves down. As soon as the supplies database is restarted, however, the rest of the services automatically bring themselves up in dependency order.

22.7 Final remarks

In this part of the thesis, we have contributed to the field of aggregate programming. The contribution can be split in three macro areas: the adaptation of classic tuple based coordinatio to space and time computation, a preliminary study on computations that are scale (namely, device-density) independent, and the introduction of a novel language, based on Field Calculus, and featuring higher order functions and integration with ALCHEMIST.

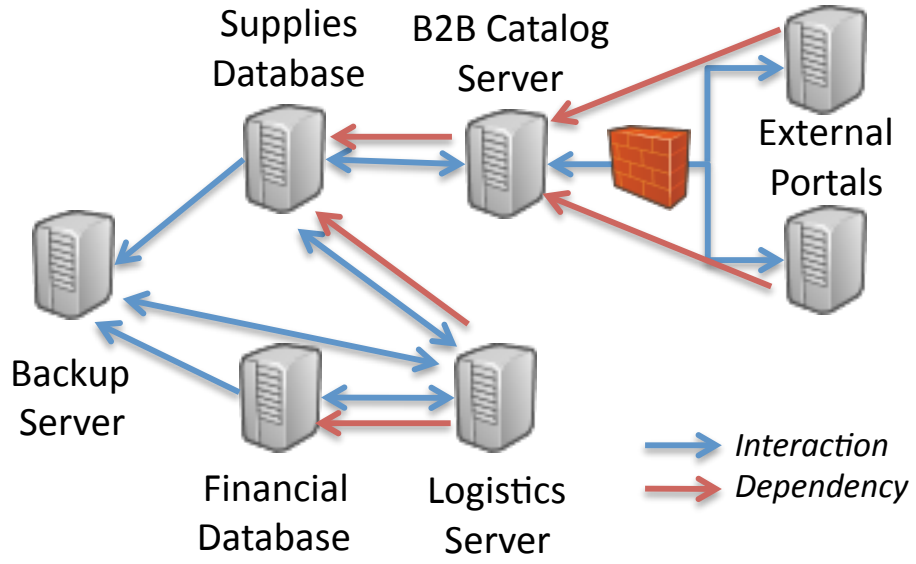

(a) Example Dependent Services Scenario

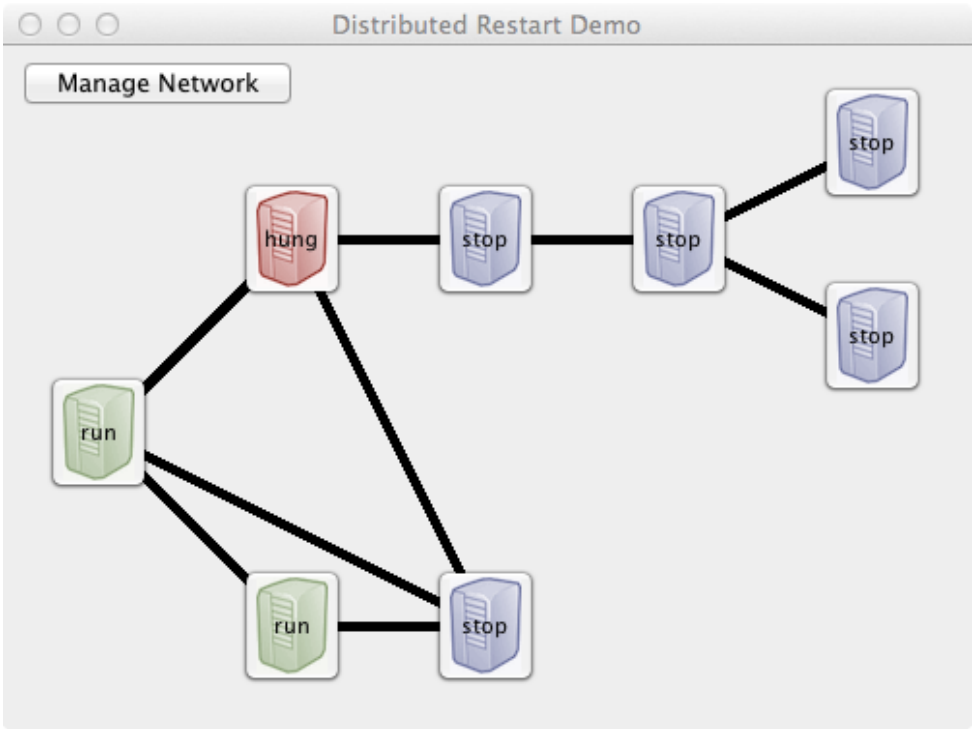

(b) Example of Coordinated Restart Execution

Figure 22.4: (a) An example scenario of an enterprise network for a small manufacturing and supply company. (b) Example of execution on a network of 8 EmuLab [WLS+02] machines: the supplies database has crashed (red), and so all dependent services have shut themselves down (blue), while other services continue to run normally (green).

Part V

Conclusion

23 Results achieved and future works

This chapter summarises the most notable contributions of this thesis.

23.1 Integrated toolchain for pervasive ecosystems

The first relevant novelty is the introduction of a toolchain for pervasive ecosystems, built from scratch taking inspiration from stochastic simulators (Chapter 4), first of its kind to the best of our knowledge.

This platform exposes a generic meta-meta model (Chapter 5) and supports a variety of concrete incarnations at the same time, making it possible to rely on such architecture for a plethora of different usages. As demonstration of flexibility, we have shown a variety of examples relying on the SAPERE incarnation (Chapter 10), one example that uses the chemical incarnation (Chapter 9) and, finally, an example that runs Protelis programs inside the platform (Section 22.5). The environments and applications varied greatly, ranging from the simulation of a morphogenesis process (Chapter 9) to crowd steering in a city, relying on existing GPS traces and realistic navigation system (Section 8.2).

ALCHEMIST itself can be classified a discrete event simulator, and its model suits well a certain category of agent based model (Chapter 7). The engine was conceived with performance in mind, and it demonstrated its qualities when compared to similar ABMs, such as Repast Symphony (Chapter 6).

The toolchain also comprises a distributed statistical model checker (MULTIVESTA), which allows the engineer to evaluate properties of models with the desired confidence and approximation, also speeding up the process by re-use and distribution of computation (Section 8.1).

23.2 Methods and patterns for self organisation

We contributed in a range of advances in the area of patterns for self-organisation.

An important contributions is a process algebra for SAPERE (Chapter 11). The key contribution of a process algebra to provide a unequivocal description of the components of the infrastructure, and also serves as an executable specification.

We devised and studied some spatial patterns: the Channel (Chapter 12), that marks devices along the path connecting a source and a destination, and the Anticipative Gradient (Chapter 13). The latter represents a contribution towards anticipative adaptation: if information about future events is available, it can be exploited in a totally distributed manner to improve the overall system performance. Also, we studied a urban sport mass event dataset in order to understand which methodologies may be suitable to obtain such data (Chapter 18).

We often focussed on the problem of crowd steering (Chapter 15) and evacuation (Chapters 14 and 19): besides the importance for safety and efficiency, this problems turns out to be rather complex to face, presenting continuously moving participants whose actions in space-time must impact the outcome of the steering system. We faced the problem at different grains, from indoor evacuation (Chapter 14) to urban-level crowd steering (Chapter 19).

We developed, relying on the SAPERE metaphors, a resource discovery system (Chapter 17). To the best of our knowledge, it is the first approach using the semantic match-degree as a mean of shaping the communications between network nodes, and incorporating extrinsic contextual features into the discovery process in such a way that resource selection emerges as an active property of the bio-inspired communication process, as contrasted to *post hoc* analysis and filtering.

Finally, we started experimenting on self-composition of services, trying to embed evolutionary algorithms in pervasive system in order to let the fittest combination of services "survive" (Chapter 16).

23.3 Aggregate programming languages

This thesis contributed to the advancements in aggregate programming in three ways:

1. A novel tuple based coordination medium is devised, that closes the gap between Linda-inspired coordination and space-time programming (Chapter 20);

2. A first study on scale-independent computations in space and time is provided, building upon Field Calculus a new operator demonstrably eventually consistent (Chapter 21);

3. The implementation of a practical programming language inspired to the Field Calculus, enriched with higher order functions: Protelis (Chapter 22).

We have developed the first practical distributed language which provides guarantees that all distributed algorithms expressed in this language are resilient against changes in the number and distribution of devices in a network. This is an important step towards a more general framework for supporting open ecosystems of pervasive wireless devices, which need to provide safe and resilient services despite running a shifting set of interacting services from many unrelated software suppliers. If it is possible to ensure that the only programs that can be expressed are those that interact well together, then it will greatly reduce the cost of providing reliable services in such environments.

When we introduced Field Calculus in Section 3.5.2, we said that it is more a theoretical than a practical framework. Building on the foundation of Field Calculus, the subsequent step is to provide a framework that can be practically leveraged to express a collection of higher-level "building block" algorithms, each a simple and generalized basis element of an "algebra" of programs with desirable resilience properties [BV14]. On top of this, higher-level library APIs can be built, enabling simple and transparent construction of robust distributed systems. Any practical implementation of Field Calculus must embed a Field Calculus interpreter within an architecture that handles the pragmatics of communication, execution, and interfacing with hardware, operating system, and other software. At the same time, it is important that this system be readily portable across both simulation environments and real networked devices. Finally, both system development and maintainability are greatly enhanced if the exact same code is used for execution in all contexts. Here is where Protelis comes into play: it represents our effort to bring Field Calculus to practice, satisfying the requirements stated above. Protelis was designed upon Field Calculus rather than higher level, more safe derivatives like GPI-Calculus (Section 21.2). The reason is that we wanted to provide a language whose underlying mechanisms may allow for any kind of program to be written, and not just "well behaving" ones.

Protelis ensures universality and coherence between aggregate specification and local execution by building atop the Field Calculus introduced by Viroli et al[VDB13]. At the same time, accessibility, portability, and ease of integration are ensured by embedding Protelis within Java. This enables Protelis programs to draw on the full breadth of available Java APIs and to readily integrate with a wide range of devices and applications, as illustrated by our examples of pervasive computing simulation and networked service management. This implementation of Protelis thus forms an important component of the toolchain necessary for practical application of aggregate programming principles and methods to address real-world problems.

23.4 Future and ongoing work

This work brought notable advancements in the field of engineering for pervasive ecosystems. As most research efforts do, however, it also shed a thin ray of light upon new lines of research that deserve to be explored. In this section, we will try to list some of the work directions that are still to be pursued, both in terms of scientific research and tool development.

23.4.1 Stabilisation of the toolchain

The ALCHEMIST toolchain for engineering pervasive systems is currently under heavy development. The project has significantly grew in size since it was conceived, both because of incremental additions to incarnations (mainly the SAPERE one) and because of novel features, above all graphical and reporting capabilities. A first step towards modularisation has been taken during the first months of development, with the adoption of Apache Maven[1]. The next step will

[1] `http://maven.apache.org/`

be the complete Maven-isation of the project, with the framework available in form of small bundles that can be composed in order to get the subset of functionality required. For instance, a much lighter version of the toolchain is sufficient to run Protelis on physical devices, since no simulation engine is required: it is pretty pointless to carry all the unused machinery, especially in case the system is going to be deployed on resource-constrained boards.

In parallel with complete Maven-isation, an improvement of the current regression test suite is scheduled. The benefits of such action should be rather clear for the reader, and represent a fundamental step in order to set up a reliable continuous building system. Up to now in fact, no fixed release scheduled has been adopted for ALCHEMIST, primarily because of the irregular development speed. A continuous building system, along with a reliable test suite, will grant that regressions are promptly discovered, roll-backed and reported (and, in turn, fixed).

23.4.2 Biochemical incarnation

We developed a very simple chemical incarnation, which was a proof of concept and also a test for us to understand how powerful the toolchain could be outside the application context it was designed for. This turned to be pretty powerful, as we demonstrated qualitatively in Section 9.1. Due to the success of this work, we decided to evolve the chemical meta model of ALCHEMIST to a biochemical meta-model, providing as first-level abstractions concept commonly modeled in context of biology and biochemistry, such as cellular bonding through junctions. This incarnation already has its domain specific language, whit a mature grammar and personalised IDE based on Eclipse[2], but whose XML generator is still in pre-alpha version. Once our effort is complete, we expect the tool to become a powerful utility for those interested in simulation of biochemical scenarios.

23.4.3 Scale independent computation

We want to generalise the results of our scale independent computations: the theoretical framework of this thesis treats only the case of stationary devices; in pervasive systems, many devices are mobile, and in practice the mechanisms used by GPI-calculus appear to perform quite well on mobile devices, so it will be important to extend the theoretical framework to handle mobile devices as well. One notable concern is that mobile devices must treat questions of Galilean vs. Einsteinian relativity [TW92], which we deferred for now.

Similarly, the theory currently only handles eventual approximation and behavior at the limit of high density networks, but these properties are most useful because in practice they tend to be indicators that an algorithm also behaves well in lower density networks and before it has finished converging. It should be possible to identify stronger properties that demarcate good performance at lower densities as well.

`GPI` is a powerful operator, but it only covers a few of the commonly used "building block" resilient distributed algorithms; another area for future extension is thus to broaden the scope

[2] `https://eclipse.org/`

of distributed applications for which resilience guarantees can be provided by extension of GPI-calculus to cover additional "building block" algorithms.

23.4.4 Distributed operating systems with Protelis

The Protelis framework continues to be actively developed: we plan to enrich it in the future by adding higher-level abstractions for aggregate programming grounded on the mechanisms discussed in Chapter 21. Also, we observe that the current situation is much like that of cooperative multitasking in the operating systems world, before the development of the current kernel-user distinction and pre-emptive multitasking that dominates the world now. We think that Protelis could be a new architecture for decentralized pre-emptive management of distributed processes, which provides these same sort of capabilities for distributed systems. In particular, we want to exploit the extension to first-class functions of Field Calculus model, which enables its alignment mechanism for pre-emptive control of the scope of a distributed computation to be used for process managament. These mechanisms enable the construction of a prototype "thread manager" for pervasive and distributed systems, including the ability to preempt a process, to kill a rogue process, or to dynamically manage the scope of a process. This point of view will be deeply analysed a paper which is an ongoing work.

It will also be important to consider how static analysis, testing, and model-checking techniques can be incorporated to extend the capabilities, especially for ensuring that semantic faults in programs can be eliminated before runtime.

Bibliography

[AAC+00] Harold Abelson, Don Allen, Daniel Coore, Chris Hanson, George Homsy, Thomas F. Knight, Jr., Radhika Nagpal, Erik Rauch, Gerald Jay Sussman, and Ron Weiss. Amorphous computing. *Communications of the ACM*, 43:74–82, 2000.

[AARC+12] Seyedreza Abdollahi, Hamed Al-Raweshidy, Erik Calvo, Luis Manuel Dahlman, Klaus David, Marian Codreanu, Panagiotis Demestichas, Markus Dillinger, Miguel Egido, Markus Gruber, Stefan Kaiser, Matti Latva-aho, Peter Merz, Xavier Mestre, Emmanouil Panaousis, Stephan Pfletschinger, Mahdi Pirmoradian, Christos Politis, Bernhard Raaf, Frank Schaich, Christoph Schmelz, Rahim Tafazolli, Antti Tölli, Nikola Vucic, and Rudolf Winkelmann. Spectrum Crunch. Technical report, Net!Works European Technology Platform, 06 2012.

[ABW06] Gianluca Antonini, Michel Bierlaire, and Mats Weber. Discrete choice models of pedestrian walking behavior. *Transportation Research Part B: Methodological*, 40(8):667 – 687, 2006.

[AGS11] Michael Adeyeye and Paul Gardner-Stephen. The village telco project: a reliable and practical wireless mesh telephony infrastructure. *EURASIP J. Wireless Comm. and Networking*, 2011:78, 2011.

[AJL+02] Bruce Alberts, Alexander Johnson, Julian Lewis, Martin Raff, Keith Roberts, and Peter Walter. *Molecular Biology of the Cell*. Garland Science Textbooks. Garland Science, Fourth edition, June 2002.

[Aka74] H. Akaike. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on*, 19(6):716–723, 1974.

[AM11] Musab AlTurki and José Meseguer. Pvesta: A parallel statistical model checking and quantitative analysis tool. In *CALCO*, 2011.

[AMS06] Gul A. Agha, José Meseguer, and Koushik Sen. PMaude: Rewrite-based specification language for probabilistic object systems. In Antonio Cerone and Herbert Wiklicky, editors, *QAPL 2005*, volume 153(2) of *ENTCS*, pages 213–239. Elsevier, 2006.

[Apa13] Tatjana Apanasevic. Factors influencing the slow rate of penetration of NFC mobile payment in western europe. In *12th International Conference on Mobile Business, ICMB 2013, Berlin, Germany, 10-13 June 2013*, page 8, 2013.

[APNF13] Bernhard Anzengruber, Danilo Pianini, Jussi Nieminen, and Alois Ferscha. Predicting social density in mass events to prevent crowd disasters. In Adam Jatowt, Ee-Peng Lim, Ying Ding, Asako Miura, Taro Tezuka, Gaël Dias, Katsumi Tanaka, Andrew J. Flanagin, and Bing Tian Dai, editors, *SocInfo*, volume 8238 of *Lecture Notes in Computer Science*, pages 206–215. Springer, 2013.

[ARGL+07] Michael P. Ashley-Rollman, Seth Copen Goldstein, Peter Lee, Todd C. Mowry, and Padmanabhan Pillai. Meld: A declarative approach to programming ensembles. In *IEEE International Conference on Intelligent Robots and Systems (IROS '07)*, pages 2794–2800, 2007.

[ASS+11] Henrik Abramowicz, Ernst-Dieter Schmidt, Lars Christoph Schmelz, Cornel Pampu, Cornelia Kappler, Mirko Schramm, Konstantinous Pentikousis, Dongming Zhou, Savo Glisic, Juan Palacios, Josep Mangues, Panagiotis Demestichas, Ralf Toenjes, Tipu Arvind Ramrekha, Grant Millar, Christos Politis, Sergi Figuerola, Dimitra Simeonidou, Alex Galis, and Rahim Tafazolli. Future Networks and Management . Technical report, Net!Works European Technology Platform, 06 2011.

[BB06] Jacob Beal and Jonathan Bachrach. Infrastructure for engineered emergence on sensor/actuator networks. *IEEE Intelligent Systems*, 21(2):10–19, 2006.

[BB07] Jonathan Bachrach and Jacob Beal. Building spatial computers. Technical Report MIT-CSAIL-TR-2007-017, MIT, March 2007.

[BBF07] Jonathan Bachrach, Jacob Beal, and Takeshi Fujiwara. Continuous space-time semantics allow adaptive program execution. In *IEEE SASO 2007*, pages 315–319, New York, July 2007. IEEE.

[BBG06] Maurice Beek, Antonio Bucchiarone, and Stefania Gnesi. A survey on service composition approaches: From industrial standards to formal methods. In *Technical Report 2006TR-15, Istituto*, pages 15–20. IEEE CS Press, 2006.

[BBL+11] Kevin Burrage, PamelaM. Burrage, André Leier, Tatiana Marquez-Lago, and Jr Nicolau, DanV. Stochastic simulation for spatial modelling of dynamic processes in a living cell. In Heinz Koeppl, Gianluca Setti, Mario di Bernardo, and Douglas Densmore, editors, *Design and Analysis of Biomolecular Circuits*, pages 43–62. Springer New York, 2011.

[BBVT08] Jacob Beal, Jonathan Bachrach, Daniel Vickery, and Mark Tobenkin. Fast self-healing gradients. In *Proceedings of ACM SAC 2008*, pages 1969–1975, 2008.

[BCD+06] Özalp Babaoglu, Geoffrey Canright, Andreas Deutsch, Gianni Di Caro, Frederick Ducatelle, Luca Maria Gambardella, Niloy Ganguly, Márk Jelasity, Roberto Montemanni, Alberto Montresor, and Tore Urnes. Design patterns from biology for distributed computing. *TAAS*, 1(1):26–66, 2006.

[BCM+03] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, New York, NY, USA, 2003.

[BDT99] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 1999.

[BDU+13] Jacob Beal, Stefan Dulman, Kyle Usbeck, Mirko Viroli, and Nikolaus Correll. Organizing the aggregate: Languages for spatial computing. In Marjan Mernik, editor, *Formal and Practical Aspects of Domain-Specific Languages: Recent Developments*, chapter 16, pages 436–501. IGI Global, 2013. A longer version available at: `http://arxiv.org/abs/1202.5509`.

[Bea05] Jacob Beal. Programming an amorphous computational medium. In Jean-Pierre Banatre, Pascal Fradet, Jean-Louis Giavitto, and Olivier Michel, editors, *Unconventional Programming Paradigms*, volume 3566 of *Lecture Notes in Computer Science*, pages 121–136. Springer Berlin Heidelberg, 2005.

[Bea09] Jacob Beal. Flexible self-healing gradients. In *Proceedings of the 2009 ACM Symposium on Applied Computing (SAC)*, pages 1197–1201. ACM, 2009.

[Bea10] Jacob Beal. A basis set of operators for space-time computations. In *Spatial Computing Workshop*, 2010. Available at: `http://www.spatial-computing.org/scw10/`.

[BEH+06] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava. Participatory sensing. In *In: Workshop on World-Sensor-Web (WSW'06): Mobile Device Centric Sensor Networks and Applications*, pages 117–134, 2006.

[BEN+04] Michel Berkelaar, Kjell Eikland, Peter Notebaert, et al. lpsolve: Open source (mixed-integer) linear programming system. *Eindhoven U. of Technology*, 2004.

[Ber92a] Gérard Berry. The chemical abstract machine. *Theoretical Computer Science*, 96(1):217–248, April 1992.

[Ber92b] Alan A. Berryman. The origins and evolution of predator-prey theory. *Ecology*, 73(5):1530–1535, October 1992.

[BFLM01] Jean-Pierre Banâtre, Pascal Fradet, and Daniel Le Métayer. Gamma and the chemical reaction model: Fifteen years after. In Cristian S. Calude, Gheorghe Păun, Grzegorz Rozenberg, and Arto Salomaa, editors, *Multiset Processing. Mathematical, Computer Science, and Molecular Computing Points of View*, volume 2235 of *LNCS*, pages 17–44. Springer, 2001.

[BGZ97] Nadia Busi, Roberto Gorrieri, and Gianluigi Zavattaro. On the turing equivalence of linda coordination primitives. *Electr. Notes Theor. Comput. Sci.*, 7:75, 1997.

[BGZ98] Nadia Busi, Roberto Gorrieri, and Gianluigi Zavattaro. A process algebraic view of linda coordination primitives. *Theor. Comput. Sci.*, 192(2):167–199, 1998.

[BKSZ01] C Burstedde, K Klauck, A Schadschneider, and J Zittartz. Simulation of pedestrian dynamics using a two-dimensional cellular automaton. *Physica A: Statistical Mechanics and its Applications*, 295(3–4):507 – 525, 2001.

[BLM90] Jean-Pierre Banâtre and Daniel Le Métayer. The GAMMA model and its discipline of programming. *Science of Computer Programming*, 15(1):55–77, November 1990.

[BLM13] Luca Bortolussi, Diego Latella, and Mieke Massink. Stochastic process algebra and stability analysis of collective systems. In *COORDINATION*, volume 7890 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2013.

[BMM02] Özalp Babaoglu, Hein Meling, and Alberto Montresor. Anthill: A framework for the development of agent-based peer-to-peer systems. In *Proceedings of the 22 Nd International Conference on Distributed Computing Systems (ICDCS'02)*, ICDCS '02, pages 15–, Washington, DC, USA, 2002. IEEE Computer Society.

[BMP+11] Chatschik Bisdikian, Bernhard Mitschang, Dino Pedreschi, Vincent S. Tseng, and Claudio Bettini. Challenges for mobile data management in the era of cloud and social computing. In *12th IEEE International Conference on Mobile Data Management*, page 6, 2011.

[BMS11] Jacob Beal, Olivier Michel, and Ulrik Pagh Schultz. Spatial computing: Distributed systems that take advantage of our geometric world. *ACM Transactions on Autonomous and Adaptive Systems*, 6:11:1–11:3, June 2011.

[BMV09a] Stefania Bandini, Sara Manzoni, and Giuseppe Vizzari. Agent based modeling and simulation: An informatics perspective. *Journal of Artificial Societies and Social Simulation*, 12:4, 2009.

[BMV09b] Stefania Bandini, Sara Manzoni, and Giuseppe Vizzari. Crowd Behavior Modeling: From Cellular Automata to Multi-Agent Systems. In Adelinde M. Uhrmacher and Danny Weyns, editors, *Multi-Agent Systems: Simulation and Applications*, Computational Analysis, Synthesis, and Design of Dynamic Systems, chapter 13, pages 389–418. CRC Press, June 2009.

[Bon99] Eric Bonabeau. Editor's introduction: Stigmergy. *Artificial Life*, 5(2):95–96, Spring 1999.

[BP09] Jean-Pierre Banâtre and Thierry Priol. Chemical programming of future service-oriented architectures. *JSW*, 4(7):738–746, 2009.

[BPd+08] A. Bandara, T. Payne, D. de Roure, N. Gibbins, and T. Lewis. Semantic resource matching for pervasive environments: The approach and its evaluation. Technical report, University of Southampton, 2008.

[Bro04] R.G. Brown. *Smoothing, Forecasting and Prediction of Discrete Time Series*. Dover Phoenix Editions. Dover Publications, 2004.

[BS11] F. Bobillo and U. Straccia. Reasoning with the finitely many-valued Łukasiewicz fuzzy Description Logic. *Information Sciences*, 181(4):758 – 778, 2011.

[BSB+03] J. Mark Bull, L. A. Smith, C. Ball, L. Pottage, and R. Freeman. Benchmarking java against c and fortran for scientific applications. *Concurrency and Computation: Practice and Experience*, 15(3-5):417–430, 2003.

[BUB13] Jacob Beal, Kyle Usbeck, and Brett Benyo. On the evaluation of space-time functions. *The Computer Journal*, 56(12):1500–1517, 2013. doi: 10.1093/comjnl/bxs099.

[But02] William Butera. *Programming a Paintable Computer*. PhD thesis, MIT, Cambridge, MA, USA, 2002.

[BV07] Stefania Bandini and Giuseppe Vizzari. Regulation function of the environment in agent-based simulation. In *Proceedings of the 3rd international conference on Environments for multi-agent systems III*, E4MAS'06, pages 157–169, Berlin, Heidelberg, 2007. Springer-Verlag.

[BV14] Jake Beal and Mirko Viroli. Building blocks for aggregate programming of self-organising applications. In *2nd FoCAS Workshop on Fundamentals of Collective Systems*, pages 1–6. IEEE CS, to appear, 2014.

[BVD14] Jacob Beal, Mirko Viroli, and Ferruccio Damiani. Towards a unified model of spatial computing. In *7th Spatial Computing Workshop (SCW 2014)*, AAMAS 2014, Paris, France, May 2014.

[BVPD14] Jacob Beal, Mirko Viroli, Danilo Pianini, and Ferruccio Damiani. Scale-independent computations in situated networks. *IEEE Transaction on Emergent Topics in Computing*, 2014. Submitted for publication.

[BW08] Eduard Babulak and Ming Wang. Discrete event simulation: State of the art. *iJOE*, 4(2):60–63, 2008.

[CDB+12] Marco Conti, Sajal K. Das, Chatschik Bisdikian, Mohan Kumar, Lionel M. Ni, Andrea Passarella, George Roussos, Gerhard Tröster, Gene Tsudik, and Franco Zambonelli. Looking ahead in pervasive computing: Challenges and opportunities in the era of cyber-physical convergence. *Pervasive and Mobile Computing*, 8(1):2–21, 2012.

[CDPEV05] Gerardo Canfora, Massimiliano Di Penta, Raffaele Esposito, and Maria Luisa Villani. An approach for qos-aware service composition based on genetic algorithms. In *Proceedings of the 2005 conference on Genetic and evolutionary computation*, GECCO '05, pages 1069–1075, New York, NY, USA, 2005. ACM.

[CG00] Luca Cardelli and Andrew D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, June 2000.

[CG09] Federica Ciocchetta and Maria Luisa Guerriero. Modelling biological compartments in bio-pepa. *Electr. Notes Theor. Comput. Sci.*, 227:77–95, 2009.

[CG10] Luca Cardelli and Philippa Gardner. Processes in space. In *Programs, Proofs, Processes, 6th Conference on Computability in Europe, CiE 2010*, volume 6158 of *Lecture Notes in Computer Science*, pages 78–87. Springer, 2010.

[Chu32] Alonzo Church. A set of postulates for the foundation of logic. *Annals of mathematics*, 33(2):346–366, 1932.

[CLMZ03] Giacomo Cabri, Letizia Leonardi, Marco Mamei, and Franco Zambonelli. Location-dependent services for mobile users. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 33(6):667–681, 2003.

[CLZ00] Giacomo Cabri, Letizia Leonardi, and Franco Zambonelli. MARS: A programmable coordination architecture for mobile agents. *IEEE Internet Computing*, 4(4):26–35, 2000.

[Cow00] Richard Cowan. Stochastic models for dna replication. In *In Handbook of Statistics*. Elsevier, 2000.

[CV95] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.

[CVG09] Matteo Casadei, Mirko Viroli, and Luca Gardelli. On the collective sort problem for distributed tuple spaces. *Science of Computer Programming*, 74(9):702–722, 2009.

[DAFMDMS14] Francesco Luca De Angelis, Jose Luis Fernandez-Marquez, and Giovanna Di Marzo Serugendo. Self-composition of services in pervasive systems: A chemical-inspired approach. In Gordan Jezic, Mario Kusek, Ignac Lovrek, Robert J. Howlett, and Lakhmi C. Jain, editors, *Agent and Multi-Agent Systems: Technologies and Applications*, volume 296 of *Advances in Intelligent Systems and Computing*, pages 37–46. Springer International Publishing, 2014.

[DAS01] A. K. Dey, G. D. Abowd, and D. Salber. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human Computer Interaction*, 16(2):97–166, December 2001.

[DDF+06] Simon Dobson, Spyros Denazis, Antonio Fernández, Dominique Gaïti, Erol Gelenbe, Fabio Massacci, Paddy Nixon, Fabrice Saffre, Nikita Schmidt, and Franco Zambonelli. A survey of autonomic communications. *ACM Transactions on Autonomous and Adaptive Systems*, 1(2):223–259, December 2006.

[DG08] Jeffrey Dean and Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[DLM05] Rocco De Nicola, Diego Latella, and Mieke Massink. Formal modeling and quantitative analysis of klaim-based mobile systems. In *SAC*, pages 428–435. ACM, 2005.

[DMSFMDA14] Giovanna Di Marzo Serugendo, Jose Luis Fernandez-Marquez, and Francesco Luca De Angelis. Engineering spatial services: Concepts, architecture, and execution models. In *Handbook of Research on Architectural Trends in Service-Driven Computing*, pages 136–159. IGI Global, 2014.

[DS04] Marco Dorigo and Thomas Stützle. *Ant Colony Optimization*. MIT Press, Cambridge, MA, July 2004.

[DWH07] T. De Wolf and T. Holvoet. Design patterns for decentralised coordination in self-organising emergent systems. *Engineering Self-Organising Systems*, 4335:28–49, February 2007.

[EAWS12] Thaddeus Eze, Richard Anthony, Chris Walshaw, and Alan Soper. Autonomic computing in the first decade: trends and direction. In *ICAS 2012, The Eighth International Conference on Autonomic and Autonomous Systems*, pages 80–85, 2012.

[EB10] Moritz Eysholdt and Heiko Behrens. Xtext: implement your language faster than the quick and dirty way. In *OOPSLA*, pages 307–309. ACM, 2010.

[EC89] Bradley R. Engstrom and Peter R. Cappello. The sdef programming system. *Journal of Parallel and Distributed Computing*, 7(2):201 – 231, 1989.

[EMRU07] Roland Ewald, Carsten Maus, Arndt Rolfs, and Adelinde M. Uhrmacher. Discrete event modeling and simulation in systems biology. *Journal of Simulation*, 1(2):81–96, 2007.

[EP05] N. Eagle and A. Pentland. Social serendipity: mobilizing social software. *Pervasive Computing, IEEE*, 4(2):28–34, 2005.

[FAS+14] Jose Luis Fernandez-Marquez, Francesco De Angelis, Giovanna Di Marzo Serugendo, Graeme Stevenson, and Gabriella Castelli. The ONE-SAPERE simulator: A prototyping tool for engineering self-organisation in pervasive environments. In *Eighth IEEE International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2014, London, United Kingdom, September 8-12, 2014*, pages 201–202. IEEE, 2014.

[FDMSS10] Regina Frei, Giovanna Di Marzo Serugendo, and Traian-Florin Serbanuta. Ambient intelligence in self-organising assembly systems using the chemical reaction model. *Journal of Ambient Intelligence and Humanized Computing*, 1(3):163–184, 2010.

[FHA99] Eric Freeman, Susanne Hupfer, and Ken Arnold. *JavaSpaces Principles, Patterns, and Practice: Principles, Patterns and Practices*. The Jini Technology Series. Addison-Wesley Longman, June 1999.

[FMDMSM12a] Jose Luis Fernandez-Marquez, Giovanna Di Marzo Serugendo, and Sara Montagna. Bio-core: Bio-inspired self-organising mechanisms core. In Emma Hart, Jon Timmis, Paul Mitchell, Takadash Nakamo, Foad Dabiri, Ozgur Akan, Paolo Bellavista, Jiannong Cao, Falko Dressler, Domenico Ferrari, Mario Gerla, Hisashi Kobayashi,

Sergio Palazzo, Sartaj Sahni, Xuemin (Sherman) Shen, Mircea Stan, Jia Xiaohua, Albert Zomaya, and Geoffrey Coulson, editors, *Bio-Inspired Models of Networks, Information, and Computing Systems*, volume 103 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 59–72. Springer Berlin Heidelberg, 2012.

[FMDMSM+12b] Jose Luis Fernandez-Marquez, Giovanna Di Marzo Serugendo, Sara Montagna, Mirko Viroli, and Josep Lluis Arcos. Description and composition of bio-inspired design patterns: a complete overview. *Natural Computing*, May 2012. Online First.

[FMST+12] Jose Luis Fernandez-Marquez, Graeme Stevenson, Akla Esso Tchao, Juan Ye, Giovanna Di Marzo Serugendo, and Simon Dobson. Analysis of new gradient based aggregation algorithms for data-propagation in distributed networks. In José Luis Fernandez-Marquez, Sara Montagna, Andrea Omicini, and Franco Zambonelli, editors, *1st International Workshop on Adaptive Service Ecosystems: Natural and Socially Inspired Solutions (ASENSIS 2012)*, SASO 2012, Lyon, France, september 2012.

[Fru93] J.J Fruin. The causes and prevention of crowd disasters. *R. A. Smith and J. F. Dickie, eds, Engineering for Crowd Safety*, pages 99–108, 1993.

[GB00] Michael A. Gibson and Jehoshua Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *J. Phys. Chem. A*, 104:1876–1889, 2000.

[Gel85] David Gelernter. Generative communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, January 1985.

[GGG05] Ramakrishna Gummadi, Omprakash Gnawali, and Ramesh Govindan. Macroprogramming wireless sensor networks using kairos. In *Distributed Computing in Sensor Systems (DCOSS)*, pages 126–140, 2005.

[GHJV95] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of reusable Object-Oriented Software*. Addison-Wesley, Reading, Mass., 1995.

[Gil77] Daniel T. Gillespie. Exact stochastic simulation of coupled crfreactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, December 1977.

[Gil06] Scott F. Gilbert. *Developmental Biology, Eighth Edition*. Sinauer Associates Inc., Eighth edition, March 2006.

[GJK+04] Vitaly V. Gursky, Johannes Jaeger, Konstantin N. Kozlov, John Reinitz, and Alexander M. Samsonov. Pattern formation and nuclear divisions are uncoupled in drosophila segmentation: comparison of spatially discrete and continuous models. *Physica D: Nonlinear Phenomena*, 197(3-4):286–302, October 2004.

[GJM12] Hongliang Guo, Yaochu Jin, and Yan Meng. A morphogenetic framework for self-organized multirobot pattern formation and boundary coverage. *ACM Transactions on Autonomous and Adaptive Systems*, 7(1):15, 2012.

[GLZ06] Roberto Gorrieri, Roberto Lucchi, and Gianluigi Zavattaro. Supporting secure coordination in secspaces. *Fundam. Inform.*, 73(4):479–506, 2006.

[GMCS05] Jean-Louis Giavitto, Olivier Michel, Julien Cohen, and Antoine Spicher. Computations in space and space in computations. In Jean-Pierre Banâtre, Pascal Fradet, Jean-Louis Giavitto, and Olivier Michel, editors, *Unconventional Programming Paradigms*, volume 3566 of *Lecture Notes in Computer Science*, pages 137–152. Springer Berlin Heidelberg, 2005.

[GPO13] Maria Ganzha, Marcin Paprzycki, and Andrea Omicini. Software agents: Twenty years and counting. *Computing Now*, 6(11), November 2013.

[GPOS13] Pedro Pablo González Pérez, Andrea Omicini, and Marco Sbaraglia. A biochemically-inspired coordination-based model for simulating intracellular signalling pathways. *Journal of Simulation*, 7(3):216–226, August 2013. Special Issue: Agent-based Modeling and Simulation.

[GPVD99] E. Guttman, C. Perkins, J. Veizades, and M. Day. Service Location Protocol, Version 2. RFC 2608, 1999.

[Gra59] Pierre-Paul Grassé. La reconstruction du nid et les coordinations interindividuelles chez Bellicositermes natalensis et Cubitermes sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes Sociaux*, 6(1):41–80, March 1959.

[GSP11] Paul Gardner-Stephen and Swapna Palaniswamy. Serval mesh software-wifi multi model management. In *Proceedings of the 1st International Conference on Wireless Technologies for Humanitarian Relief*, ACWR '11, pages 71–77, New York, NY, USA, 2011. ACM.

[GSW06] Dina Q. Goldin, Scott A. Smolka, and Peter Wegner, editors. *Interactive Computation: The New Paradigm*. Springer, September 2006.

[GVO07] Luca Gardelli, Mirko Viroli, and Andrea Omicini. Design patterns for self-organising systems. In Hans-Dieter Burkhard, Rineke Verbrugge, and László Zsolt Varga, editors, *Multi-Agent Systems and Applications V*, volume 4696 of *LNAI*, pages 123–132. Springer, September 2007. 5th International Central and Eastern European Conference on Multi-Agent Systems (CEEMAS'07), Leipzig, Germany, 25–27 September 2007, Proceedings.

[HAES09] M. Shamim Hossain, Atif Alamri, and Abdulmotaleb El Saddik. A biologically inspired framework for multimedia service management in a ubiquitous environment. *Concurrency and Computation: Practice and Experience*, 21(11):1450–1466, 2009.

[Har87] David Harel. Statecharts: A visual formalism for complex systems. *Sci. Comput. Program.*, 8(3):231–274, June 1987.

[HBJW05] Dirk Helbing, Lubos Buzna, Anders Johansson, and Torsten Werner. Self-organized pedestrian crowd dynamics: Experiments, simulations, and design solutions. *Transportation Science*, 39(1):1–24, 2005.

[HDBDM12] Dries Harnie, Theo D'Hondt, Elisa Gonzalez Boix, and Wolfgang De Meuter. Programming urban-area applications. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing*, pages 1516–1521. ACM, 2012.

[HDF+11] Philip N. Howard, Aiden Duffy, Deen Freelon, Muzammil Hussain, Will Mari, and Marwa Mazaid. Opening closed regimes: What was the role of social media during the arab spring? 2011.

[Hen96] T.A. Henzinger. The theory of hybrid automata. In *Logic in Computer Science, 1996. LICS '96. Proceedings., Eleventh Annual IEEE Symposium on*, pages 278–292, Jul 1996.

[HFH+09] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. The weka data mining software: an update. *SIGKDD Explor. Newsl.*, 11(1):10–18, November 2009.

[HJ09] D. Helbing and A. Johansson. *Pedestrian, crowd, and evacuation dynamics.*, volume 16, pages 6476–6495. Springer, New York, 2009.

[HLT11] Wei-Lieh Hsu, Kun-Fong Lin, and Chang-Lung Tsai. Crowd density estimation based on frequency analysis. In *Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP), 2011 Seventh International Conference on*, pages 348–351, 2011.

[HM12] Dirk Helbing and Pratik Mukerji. Crowd disasters as systemic failures: analysis of the love parade disaster. *EPJ Data Science*, 1(1):7, 2012.

[HQL+11] Fei Hu, Meikang Qiu, Jiayin Li, Travis Grant, Drew Taylor, Seth McCaleb, Lee Butler, and Richard Hamner. A review on cloud computing: Design challenges in architecture and security. *CIT. Journal of Computing and Information Technology*, 19(1):25–55, 2011.

[HSG+06] Stefan Hoops, Sven Sahle, Ralph Gauges, Christine Lee, Jörgen Pahle, Natalia Simus, Mudita Singhal, Liang Xu, Pedro Mendes, and Ursula Kummer. Copasi – a complex pathway simulator. *Bioinformatics*, 22(24):3067–3074, 2006.

[HVUM07] Alexander Helleboogh, Giuseppe Vizzari, Adelinde Uhrmacher, and Fabien Michel. Modeling dynamic environments in multi-agent simulation. *Autonomous Agents and Multi-Agent Systems*, 14(1):87–116, February 2007.

[HW08] Mordechai (Muki) Haklay and Patrick Weber. Openstreetmap: User-generated street maps. *IEEE Pervasive Computing*, 7(4):12–18, October 2008.

[IPW01] Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. Featherweight java: A minimal core calculus for Java and GJ. *ACM Transactions on Programming Languages and Systems*, 23(3), 2001.

[Jaz05] Mehdi Jazayeri. Species evolve, individuals age. In Motoshi Saeki, Gerardo Canfora, and Shuichiro Yamamoto, editors, *8th International Workshop on Principles of Software Evolution (IWPSE 2005)*, pages 3–9, Lisbon, Portugal, 5–6 September 2005.

[JBH+12] Anders Johansson, Michael Batty, Konrad Hayashi, Osama Al Bar, David Marcozzi, and Ziad A Memish. Crowd and environmental management during mass gatherings. *The Lancet Infectious Diseases*, 12(2):150 – 156, 2012.

[JHP+03] J.G. Jetcheva, Y.-C. Hu, S. PalChaudhuri, A.K. Saha, and D.B. Johnson. Design and evaluation of a metropolitan area multitier wireless ad hoc network architecture. In *Proceedings of the 5th IEEE Workshop on Mobile Computing Systems and Applications*, pages 32–43, Monterey, CA, USA, October 2003.

[KC09] Rick Kazman and Hong-Mei Chen. The metropolis model a new logic for development of crowdsourced systems. *Commun. ACM*, 52:76–84, July 2009.

[Kie02] Andrzej M. Kierzek. Stocks: Stochastic kinetic simulations of biochemical systems with gillespie algorithm. *Bioinformatics*, 18(3):470–481, 2002.

[KR08] Lila Kari and Grzegorz Rozenberg. The many facets of natural computing. *Communications of the ACM*, 51:72–83, 2008.

[KRS10] W.W.F. Klingsch, C. Rogsch, and A. Schadschneider. *Pedestrian and Evacuation Dynamics 2008*. Springer Berlin Heidelberg, 2010.

[KSBM01] S. S. Keerthi, S. K. Shevade, C. Bhattacharyya, and K. R. K. Murthy. Improvements to platt's smo algorithm for svm classifier design. *Neural Comput.*, 13(3):637–649, March 2001.

[LBF13] Ivan Lanese, Luca Bedogni, and Marco Di Felice. Internet of things: a process calculus approach. In *SAC*, pages 1339–1346. ACM, 2013.

[LCRP+05] Sean Luke, Claudio Cioffi-Revilla, Liviu Panait, Keith Sullivan, and Gabriel Catalin Balan. Mason: A multiagent simulation environment. *Simulation*, 81(7):517–527, 2005.

[LDB09] William J.R. Longabaugh, Eric H. Davidson, and Hamid Bolouri. Visualization, documentation, analysis, and communication of large-scale gene regulatory networks. *Biochimica et Biophysica Acta (BBA) - Gene Regulatory Mechanisms*, 1789(4):363–374, 2009.

[LH03] L. Li and I. Horrocks. A software framework for matchmaking based on Semantic Web technology. In *Proceedings of the 12th international conference on World Wide Web*, pages 331–339, New York, NY, USA, 2003.

[LHK+09] Kyunghan Lee, Seongik Hong, Seong Joon Kim, Injong Rhee, and Song Chong. Slaw: A mobility model for human walks. In *Proceedings of the 28th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Rio de Janeiro, Brazil, April 2009. IEEE.

[LIDP10] Paola Lecca, Adaoha E. C. Ihekwaba, Lorenzo Dematté, and Corrado Priami. Stochastic simulation of the spatio-temporal dynamics of reaction-diffusion systems: the case for the bicoid gradient. *J. Integrative Bioinformatics*, 7(1), 2010.

[LK87] Frank C. H. Lin and Robert M. Keller. The gradient model load balancing method. *IEEE Trans. Softw. Eng.*, 13(1):32–38, 1987.

[LMG+09] Uichin Lee, Eugenio Magistretti, Mario Gerla, Paolo Bellavista, Pietro Lió, and Kang-Won Lee. Bio-inspired multi-agent data harvesting in a proactive urban monitoring environment. *Ad Hoc Networks*, 7:725–741, 2009.

[LMMD88] C. Lasser, J.P. Massar, J. Miney, and L. Dayton. *Starlisp Reference Manual*. Thinking Machines Corporation, 1988.

[LPF12] P. Lukowicz, S. Pentland, and A. Ferscha. From context awareness to socially aware computing. *IEEE Pervasive Computing*, 11(1):32 –41, 2012.

[LT06] Jiming Liu and Kwok Ching Tsui. Toward nature-inspired computing. *Communications of the ACM*, 49(10):59–64, October 2006.

[Lyn96] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1996.

[Mac90] Bruce MacLennan. Continuous spatial automata. Technical Report Department of Computer Science Technical Report CS-90-121, University of Tennessee, Knoxville, November 1990.

[Mac98] David J.C. Mackay. Introduction to gaussian processes, 1998.

[Mac08] Rebecca MacKinnon. Flatter world and thicker walls? blogs, censorship and civic discourse in china. *Public Choice*, 134(1-2):31–46, 2008.

[MBGH07] Kaj Mäkelä, Sara Belt, Dan Greenblatt, and Jonna Häkkilä. Mobile interaction with visual and rfid tags: a field study on user perceptions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 991–994, New York, NY, USA, 2007. ACM.

[MCOV11] Ambra Molesini, Matteo Casadei, Andrea Omicini, and Mirko Viroli. Simulation in agent-oriented software engineering: The SODA case study. *Science of Computer Programming*, August 2011. Special Issue on Agent-oriented Design methods and Programming Techniques for Distributed Computing in Dynamic and Complex Environments.

[MFHH05] S. Madden, M. Franklin, J. Hellerstein, and W. Hong. Tinydb: An acqusitional query processing system for sensor networks. In *ACM TODS*, 2005.

[Mil99] Robin Milner. *Communicating and Mobile Systems: The Pi-Calculus*. Cambridge University Press, 1999.

[ML12] Mieke Massink and Diego Latella. Fluid analysis of foraging ants. In *COORDINATION*, volume 7274 of *Lecture Notes in Computer Science*, pages 152–165. Springer, 2012.

[MLHT04] R. Ma, Liyuan Li, Weimin Huang, and Qi Tian. On pixel count based crowd density estimation for visual surveillance. In *Cybernetics and Intelligent Systems, 2004 IEEE Conference on*, volume 1, pages 170–173 vol.1, 2004.

[MM04] Frank Manola and Eric Miller. RDF primer. W3C recommendation, World Wide Web Consortium, February 2004.

[MMTL13] Massimo Monti, Thomas Meyer, Christian F. Tschudin, and Marco Luise. Stability and sensitivity analysis of traffic-shaping algorithms inspired by chemical engineering. *IEEE Journal on Selected Areas in Communications*, 31(6):1105–1114, 2013.

[MMTZ06] M. Mamei, R. Menezes, R. Tolksdorf, and F. Zambonelli. Case studies for self-organization in computer science. *Journal of Systems Architecture*, 52:443–460, 2006.

[MN98] Makoto Matsumoto and Takuji Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.*, 8(1):3–30, 1998.

[MN10] Charles M. Macal and Michael J. North. Tutorial on agent-based modelling and simulation. *Journal of Simulation*, 4(3):151–162, 2010.

[MO13] Stefano Mariani and Andrea Omicini. Molecules of knowledge: Self-organisation in knowledge-intensive environments. *Intelligent Distributed Computing VI*, pages 17–22, 2013.

[Mon81] Gaspard Monge. Memoire sur la theorie des delais et des remblais. *Histoire de l'Academie Royale des Sciences de Paris, avec les Memoires de Mathematique et de Physique pour la meme annee*, pages 666–704, 1781.

[MOV09] Ambra Molesini, Andrea Omicini, and Mirko Viroli. Environment in Agent-Oriented Software Engineering methodologies. *Multiagent and Grid Systems*, 5(1):37–57, 2009. Special Issue "Engineering Environments in Multi-Agent Systems".

[MPR06] Amy L. Murphy, Gian Pietro Picco, and Gruia-Catalin Roman. Lime: A model and middleware supporting mobility of hosts and agents. *ACM Trans. on Software Engineering and Methodology*, 15(3):279–328, 2006.

[MPV12a] Sara Montagna, Danilo Pianini, and Mirko Viroli. Gradient-based self-organisation patterns of anticipative adaptation. In *Proceedings of 6th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2012)*, pages 169–174, September 2012.

[MPV12b] Sara Montagna, Danilo Pianini, and Mirko Viroli. A model for drosophila melanogaster development from a single cell to stripe pattern formation. In *Proceedings of the ACM Symposium on Applied Computing, SAC 2012, Riva, Trento, Italy, March 26-30, 2012*, pages 1406–1412, 2012.

[MT03] Ronaldo Menezes and Robert Tolksdorf. Adaptiveness in Linda-based coordination models. In Giovanna Di Marzo Serugendo, Anthony Karageorgos, Omer F. Rana, and Franco Zambonelli, editors, *Engineering Self-Organising Systems. Nature-Inspired Approaches to Software Engineering*, volume 2977, pages 212–232. Springer, 2003.

[MTUG09] Aneil Mallavarapu, Matthew Thomson, Benjamin Ullian, and Jeremy Gunawardena. Programming with models: modularity and abstraction provide powerful capabilities for systems biology. *Journal of the Royal Society, Interface / the Royal Society*, 6(32):257–270, March 2009.

[Mur89] T. Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, April 1989.

[MV10] Sara Montagna and Mirko Viroli. A framework for modelling and simulating networks of cells. *Electronic Notes in Theoretical Computer Science*, 268:115–129, December 2010. Proceedings of the 1st International Workshop on Interactions between Computer Science and Biology (CS2Bio'10).

[MVFM+13] Sara Montagna, Mirko Viroli, Jose Luis Fernandez-Marquez, Giovanna Di Marzo Serugendo, and Franco Zambonelli. Injecting self-organisation into pervasive service ecosystems. *Mobile Networks and Applications*, 18(3):398–412, 2013.

[MVPFM12] Sara Montagna, Mirko Viroli, Danilo Pianini, and Jose Luis Fernandez-Marquez. Towards a comprehensive approach to spontaneous self-composition in pervasive ecosystems. In Flavio De Paoli and Giuseppe Vizzari, editors, *Proceedings of the 13th Workshop on Objects and Agents*. CEUR-WS, 12 September 2012.

[MVR+11] Sara Montagna, Mirko Viroli, Matteo Risoldi, Danilo Pianini, and Giovanna Di Marzo Serugendo. Self-organising pervasive ecosystems: A crowd evacuation example. In *Software Engineering for Resilient Systems - Third International Workshop, SERENE 2011, Geneva, Switzerland, September 29-30, 2011. Proceedings*, pages 115–129, 2011.

[MYT07] Thomas Meyer, Lidia Yamamoto, and Christian F. Tschudin. An artificial chemistry for networking. In *First Workshop on Bio-Inspired Design of Networks, BIOWIRE 2007, Cambridge, UK*, pages 45–57, 2007.

[MZ05] Marco Mamei and Franco Zambonelli. Programming stigmergic coordination with the TOTA middleware. In *4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005)*, pages 415–422, New York, NY, USA, 2005. ACM.

[MZ06a] Marco Mamei and Franco Zambonelli. *Field-Based Coordination for Pervasive Multiagent Systems. Models, Technologies, and Applications*. Springer Series in Agent Technology. Springer, March 2006.

[MZ06b] Marco Mamei and Franco Zambonelli. Self-maintained distributed tuples for field-based coordination in dynamic networks. *Concurrency and Computation: Practice and Experience*, 18(4):427–443, 2006.

[MZ07] Marco Mamei and Franco Zambonelli. Pervasive pheromone-based interaction with RFID tags. *TAAS*, 2(2), 2007.

[MZ09] Marco Mamei and Franco Zambonelli. Programming pervasive and mobile computing applications: The tota approach. *ACM Trans. Softw. Eng. Methodol.*, 18(4):1–56, 2009.

[MZL03] Marco Mamei, Franco Zambonelli, and Letizia Leonardi. Co-fields: Towards a unifying approach to the engineering of swarm intelligent systems. In Paolo Petta, Robert Tolksdorf, and Franco Zambonelli, editors, *Engineering Societies in the Agents World III*, volume 2577 of *LNCS*, pages 68–81. Springer, April 2003. 3rd International Workshop (ESAW 2002), Madrid, Spain, 16–17 September 2002. Revised Papers.

[Nag01] Radhika Nagpal. *Programmable Self-Assembly: Constructing Global Shape using Biologically-inspired Local Interactions and Origami Mathematics*. PhD thesis, MIT, Cambridge, MA, USA, 2001.

[NFP98] Rocco De Nicola, Gian Luigi Ferrari, and Rosario Pugliese. Klaim: A kernel language for agents interaction and mobility. *IEEE Trans. Software Eng.*, 24(5):315–330, 1998.

[NHCV07] Michael J. North, Tom R. Howe, Nicholson T. Collier, and Jerry R. Vos. A declarative model assembly infrastructure for verification and validation. In Shingo Takahashi, David Sallach, and Juliette Rouchier, editors, *Advancing Social Simulation: The First World Congress*, pages 129–140. Springer Japan, 2007.

[NIK08] M. Nakatsuka, H. Iwatani, and J. Katto. A study on passive crowd density estimation using wireless sensors. In *4th International Conference on Mobile Computing and Ubiquitous Networking (ICMU), Tokyo, Japan*, 2008.

[NLPT14] Rocco De Nicola, Michele Loreti, Rosario Pugliese, and Francesco Tiezzi. A formal approach to autonomic systems programming: The scel language. *ACM Trans. Auton. Adapt. Syst.*, 9(2):7:1–7:29, July 2014.

[NMS07] Anirudh Natarajan, Mehul Motani, and Vikram Srinivasan. Understanding urban interactions from bluetooth phone contact traces. In Steve Uhlig, Konstantina Papagiannaki, and Olivier Bonaventure, editors, *Passive and Active Network Measurement*, volume 4427 of *Lecture Notes in Computer Science*, pages 115–124. Springer Berlin Heidelberg, 2007.

[NR95] C.E. Nicholson and B. Roebuck. The investigation of the hillsborough disaster by the health and safety executive. *Safety Science*, 18(4):249 – 259, 1995.

[NW04] Ryan Newton and Matt Welsh. Region streams: Functional macroprogramming for sensor networks. In *First International Workshop on Data Management for Sensor Networks (DMSN)*, pages 78–87, August 2004.

[NZ01] Michael S. Noble and Stoyanka Zlateva. Scientific computation with javaspaces. In *Proceedings of the 9th International Conference on High-Performance Computing and Networking*, pages 657–666. Springer-Verlag, 2001.

[OC13] Andrea Omicini and Pierluigi Contucci. Complexity and interaction: Blurring borders between physical, computational, and social systems. Preliminary notes. In Costin Bădică, Ngoc Thanh Nguyen, and Marius Brezovan, editors, *Computational Collective Intelligence. Technologies and Applications*, volume 8083 of *Lecture Notes in Computer Science*, pages 1–10. Springer Berlin Heidelberg, 2013. 5th International Conference (ICCCI 2013). Craiova, Romania, 11–13 September 2013, Proceedings. Invited Paper.

[OP07] Jan Ondrus and Yves Pigneur. An assessment of nfc for future mobile payment systems. In *Proceedings of the International Conference on the Management of Mobile Business*, ICMB '07, pages 43–, Washington, DC, USA, 2007. IEEE Computer Society.

[ORAI11] Bogdan Oancea, Ion Gh. Rosca, Tudorel Andrei, and Andreea Iluzia Iacob. Evaluating java performance for linear algebra numerical computations. *Procedia CS*, 3:474–478, 2011.

[ORV05] Andrea Omicini, Alessandro Ricci, and Mirko Viroli. An algebraic approach for modelling organisation, roles and contexts in MAS. *Applicable Algebra in Engineering, Communication and Computing*, 16(2-3):151–178, August 2005.

[ORV06] Andrea Omicini, Alessandro Ricci, and Mirko Viroli. The multidisciplinary patterns of interaction from sciences to Computer Science. In *Interactive Computation: The New Paradigm*, pages 395–414. Springer, September 2006.

[ORV08] Andrea Omicini, Alessandro Ricci, and Mirko Viroli. Artifacts in the A&A meta-model for multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 17(3):432–456, December 2008. Special Issue on Foundations, Advanced Topics and Industrial Perspectives of Multi-Agent Systems.

[OZ99] Andrea Omicini and Franco Zambonelli. Coordination for Internet application development. *Autonomous Agents and Multi-Agent Systems*, 2(3):251–269, September 1999.

[Par06] H. Van Dyke Parunak. A survey of environments and mechanisms for human-human stigmergy. In Danny Weyns, H. Van Dyke Parunak, and Fabien Michel, editors, *Environments for Multi-Agent Systems II*, volume 3830 of *LNCS*, pages 163–186. Springer, 2006.

[PBS02] H. Van Dyke Parunak, Sven Brueckner, and John Sauter. Digital pheromone mechanisms for coordination of unmanned vehicles. In Cristiano Castelfranchi and W. Lewis Johnson, editors, *1st International Joint Conference on Autonomous Agents and Multiagent systems*, volume 1, pages 449–450, New York, NY, USA, 15–19July 2002. ACM.

[PBV15] Danilo Pianini, Jacob Beal, and Mirko Viroli. Practical aggregate programming with PROTELIS. In *ACM Symposium on Applied Computing (SAC 2015)*, 2015. To appear.

[PCBB07] Julien Pauty, Paul Couderc, Michel Banatre, and Yolande Berbers. Geo-linda: a geometry aware distributed tuple space. In *IEEE 21st International Conference on Advanced Networking and Applications (AINA '07)*, pages 370–377, May 2007.

[PD12] Anna-Kaisa Pietiläinen and Christophe Diot. Dissemination in opportunistic social networks: the role of temporal communities. In *Proceedings of the thirteenth ACM international symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '12, pages 165–174, New York, NY, USA, 2012. ACM.

[PH14] T.-T.T. Pham and J.C. Ho. What are the core drivers in consumer adoption of nfc-based mobile payments?: A proposed research framework. In *Management of Engineering Technology (PICMET), 2014 Portland International Conference on*, pages 3041–3049, July 2014.

[PJRG06] Theodore J Perkins, Johannes Jaeger, John Reinitz, and Leon Glass. Reverse engineering the gap gene network of *Drosophila Melanogaster*. *PLoS Computational Biology*, 2(5):e51, 05 2006.

[PKPS02] M. Paolucci, T. Kawamura, T. R. Payne, and K. P. Sycara. Semantic matching of web services capabilities. In *Proceedings of the First International Semantic Web Conference on The Semantic Web*, ISWC '02, pages 333–347, London, UK, 2002. Springer.

[PMV11] Danilo Pianini, Sara Montagna, and Mirko Viroli. A chemical inspired simulation framework for pervasive services ecosystems. In *Proceedings of the Federated Conference on Computer Science and Information Systems*, pages 667–674. IEEE Computer Society Press, 2011.

[PMV13] Danilo Pianini, Sara Montagna, and Mirko Viroli. Chemical-oriented simulation of computational systems with Alchemist. *Journal of Simulation*, 2013.

[Pol89] Lissa F. Pollacia. A survey of discrete event simulation and state-of-the-art discrete event languages. *SIGSIM Simul. Dig.*, 20(3):8–25, September 1989.

[PPSR09] Andrei Pisarev, Ekaterina Poustelnikova, Maria Samsonova, and John Reinitz. Flyex, the quantitative atlas on segmentation gene expression at cellular resolution. *Nucleic Acids Research*, 37(Database-Issue):560–566, 2009.

[Pri95] Corrado Priami. Stochastic pi-calculus. *Comput. J.*, 38(7):578–589, 1995.

[PRSS01] Corrado Priami, Aviv Regev, Ehud Shapiro, and William Silverman. Application of a stochastic name-passing calculus to representation and simulation of molecular processes. *Information Processing Letters*, 80:25–31, October 2001.

[PSV14] Danilo Pianini, Stefano Sebastio, and Andrea Vandin. Distributed statistical analysis of complex systems modeled through a chemical metaphor. In *International Conference on High Performance Computing & Simulation, HPCS 2014, Bologna, Italy, 21-25 July, 2014*, pages 416–423, 2014.

[PVM+10] Danilo Pianini, Sascia Virruso, Ronaldo Menezes, Andrea Omicini, and Mirko Viroli. Self organization in coordination systems using a wordnet-based ontology. In *SASO*, pages 114–123, 2010.

[PVM11] Danilo Pianini, Mirko Viroli, and Sara Montagna. A simulation framework for pervasive services ecosystems. In *Proceedings of the 12th Workshop on Objects and Agents, Rende (CS), Italy, Jul 4-6, 2011*, pages 150–157, 2011.

[PVZF14] Danilo Pianini, Mirko Viroli, Franco Zambonelli, and Alois Ferscha. HPC from a self-organisation perspective: The case of crowd steering at the urban scale. In *International Conference on High Performance Computing & Simulation, HPCS 2014, Bologna, Italy, 21-25 July, 2014*, pages 460–467, 2014.

[Ras04] CarlEdward Rasmussen. Gaussian processes in machine learning. In Olivier Bousquet, Ulrike Luxburg, and Gunnar Rätsch, editors, *Advanced Lectures on Machine Learning*, volume 3176 of *Lecture Notes in Computer Science*, pages 63–71. Springer Berlin Heidelberg, 2004.

[RCD01] Davide Rossi, Giacomo Cabri, and Enrico Denti. Tuple-based technologies for coordination. In Andrea Omicini, Franco Zambonelli, Matthias Klusch, and Robert Tolksdorf, editors, *Coordination of Internet Agents: Models, Technologies, and Applications*, chapter 4, pages 83–109. Springer-Verlag, March 2001.

[RCKZ13] Vaskar Raychoudhury, Jiannong Cao, Mohan Kumar, and Daqiang Zhang. Middleware for pervasive computing: A survey. *Pervasive and Mobile Computing*, 9(2):177 – 200, 2013.

[RL93] F Raimbault and D Lavenier. Relacs for systolic programming. In *Int'l Conf. on Application-Specific Array Processors*, pages 132–135, October 1993.

[RMCM04] Anand Ranganathan, Robert E. McGrath, Roy H. Campbell, and M. Dennis Mickunas. Use of ontologies in a pervasive computing environment. *Knowledge Engineering Review*, 18(3):209–220, 2004.

[ROD02] Alessandro Ricci, Andrea Omicini, and Enrico Denti. Objective vs. subjective coordination in agent-based systems: A case study. In Farhad Arbab and Carolyn L. Talcott, editors, *COORDINATION*, volume 2315 of *Lecture Notes in Computer Science*, pages 291–299. Springer, 2002.

[RPJ96] Rolando Rivera-Pomar and Herbert Jackle. From gradients to stripes in drosophila embryogenesis: filling in the gaps. *Trends in Genetics*, 12(11):478 – 483, 1996.

[RPV11] Alessandro Ricci, Michele Piunti, and Mirko Viroli. Environment programming in multi-agent systems – an artifact-based perspective. *Autonomous Agents and Multi-Agent Systems*, 23(2):158–192, September 2011.

[RS95] John Reinitz and David H. Sharp. Mechanism of eve stripe formation. *Mechanisms of Development*, 49(1-2):133–158, January 1995.

[RVO04] Alessandro Ricci, Mirko Viroli, and Andrea Omicini. Agent coordination context: From theory to practice. In Robert Trappl, editor, *Cybernetics and Systems 2004*, volume 2, pages 618–623, Vienna, Austria, 2004. Austrian Society for Cybernetic Studies. 17th European Meeting on Cybernetics and Systems Research (EMCSR 2004), Vienna, Austria, 13–16 April 2004. Proceedings.

[SCM+06] Jing Su, Kelvin K. W. Chan, Andrew G. Miklas, Kenneth Po, Ali Akhavan, Stefan Saroiu, Eyal de Lara, and Ashvin Goel. A preliminary investigation of worm infections in a bluetooth environment. In *Proceedings of the 4th ACM workshop on Recurring malcode*, WORM '06, pages 9–16, New York, NY, USA, 2006. ACM.

[SGJ07] Michael Schumacher, Laurent Grangier, and Radu Jurca. Governing environments for agent-based traffic simulations. In *Proceedings of the 5th international Central and Eastern European conference on Multi-Agent Systems and Applications V*, CEEMAS '07, pages 163–172, Berlin, Heidelberg, 2007. Springer-Verlag.

[Sha04] Nigel Shadbolt. From the Editor in Chief: Nature-Inspired Computing. *IEEE Intelligent Systems*, 19(1):2–3, January–February 2004.

[SJ07] Drew Stovall and Christine Julien. Resource discovery with evolving tuples. In *International Workshop on Engineering of Software Services for Pervasive Environments: In Conjunction with the 6th ESEC/FSE Joint Meeting*, ESSPE '07, pages 1–10, New York, NY, USA, 2007. ACM.

[SKK+08] Svetlana Surkova, David Kosman, Konstantin Kozlov, Manu, Ekaterina Myasnikova, Anastasia A. Samsonova, Alexander Spirov, Carlos E. Vanario-Alonso, Maria Samsonova, and John Reinitz. Characterization of the Drosophila segment determination morphome. *Developmental Biology*, 313(2):844 – 862, 2008.

[Skl07] Elizabeth Sklar. Netlogo, a multi-agent simulation environment. *Artificial Life*, 13(3):303–311, 2007.

[SMA05] Koushik Sen, Viswanathan M., and Gul Agha. Vesta: A statistical model-checker and analyzer for probabilistic systems. In *Quantitative Evaluation of Systems, 2005.*, pages 251–252, 2005.

[SRAA10] Norman Salazar, Juan A. Rodriguez-Aguilar, and Josep L. Arcos. Robust coordination in large convention spaces. *AI Communications*, 23(4):357–372, 2010.

[SRM+13] Daniel Schuster, Alberto Rosi, Marco Mamei, Thomas Springer, Markus Endler, and Franco Zambonelli. Pervasive social context: Taxonomy and survey. *ACM Transactions on Intelligent Systems and Technology*, 4(3), 2013.

[SRS10] Anu Singh, C. R. Ramakrishnan, and Scott A. Smolka. A process calculus for mobile ad hoc networks. *Sci. Comput. Program.*, 75(6):440–469, 2010.

[SS04] Alex J. Smola and Bernhard Schölkopf. A tutorial on support vector regression. *Statistics and Computing*, 14(3):199–222, August 2004.

[SSS07] D. Skoutas, A. Simitsis, and T. K. Sellis. A ranking mechanism for Semantic Web service discovery. In *IEEE SCW*, pages 41–48, 2007.

[STP08] Alexander Slepoy, Aidan P. Thompson, and Steven J. Plimpton. A constant-time kinetic monte carlo algorithm for simulation of large biochemical reaction networks. *The Journal of Chemical Physics*, 128(20):205101, 2008.

[SV13] Stefano Sebastio and Andrea Vandin. MultiVeStA: Statistical model checking for discrete event simulators. In *VALUETOOLS*, 2013.

[SYD+13] Graeme Stevenson, Juan Ye, Simon Dobson, Danilo Pianini, Sara Montagna, and Mirko Viroli. Combining self-organisation, context-awareness and semantic reasoning: the case of resource discovery in opportunistic networks. In Sung Y. Shin and José Carlos Maldonado, editors, *SAC*, pages 1369–1376, Coimbra, Portugal, 18–22 March 2013. ACM.

[TB99] Guy Theraulaz and Eric Bonabeau. A brief history of stigmergy. *Artificial Life*, 5(2):97–116, Spring 1999.

[TM03] Robert Tolksdorf and Ronaldo Menezes. Using swarm intelligence in linda systems. In Andrea Omicini, Paolo Petta, and Jeremy Pitt, editors, *ESAW*, volume 3071 of *Lecture Notes in Computer Science*, pages 49–65. Springer, 2003.

[TNN12] Michal Terepeta, Hanne Riis Nielson, and Flemming Nielson. Recursive advice for coordination. In *Coordination Models and Languages - 14th International Conference, COORDINATION 2012, Stockholm, Sweden, June 14-15, 2012. Proceedings*, volume 7274 of *Lecture Notes in Computer Science*, pages 137–151. Springer, 2012.

[TOCH14] Garry Wei-Han Tan, Keng-Boon Ooi, Siong-Choy Chong, and Teck-Soon Hew. Nfc mobile credit card: The next frontier of mobile payment? *Telemat. Inf.*, 31(2):292–307, May 2014.

[TW92] Edwin F. Taylor and John Archibald Wheeler. *Spacetime Physics: Introduction to Special Relativity*. W. H. Freeman & Company, 2nd ed. edition, 1992.

[UP05] Adelinde M. Uhrmacher and Corrado Priami. Discrete event systems specification in systems biology - a discussion of stochastic pi calculus and devs. In *Proceedings of the 37th conference on Winter simulation*, WSC '05, pages 317–326. Winter Simulation Conference, 2005.

[VB08] Cristian Versari and Nadia Busi. Efficient stochastic simulation of biological systems with multiple variable volumes. *Electronic Notes in Theoretical Computer Science*, 194(3):165–180, 2008.

[VBC11] Mirko Viroli, Jake Beal, and Matteo Casadei. Core operational semantics of Proto. In *26th Annual ACM Symposium on Applied Computing (SAC 2011)*, Tunghai University, TaiChung, Taiwan, 21–25 March 2011. ACM.

[VBU13] Mirko Viroli, Jacob Beal, and Kyle Usbeck. Operational semantics of Proto. *Science of Computer Programming*, 78(6):633–656, June 2013.

[VC09] Mirko Viroli and Matteo Casadei. Biochemical tuple spaces for self-organising coordination. In John Field and Vasco Thudichum Vasconcelos, editors, *COORDINATION*, volume 5521 of *Lecture Notes in Computer Science*, pages 143–162. Springer, 2009.

[VCMZ11] Mirko Viroli, Matteo Casadei, Sara Montagna, and Franco Zambonelli. Spatial coordination of pervasive services through chemical-inspired tuple spaces. *ACM Transactions on Autonomous and Adaptive Systems*, 6(2):14:1 – 14:24, June 2011.

[VCNO10] Mirko Viroli, Matteo Casadei, Elena Nardini, and Andrea Omicini. Towards a chemical-inspired infrastructure for self-* pervasive applications. In Danny Weyns, Sam Malek, Rogério de Lemos, and Jesper Andersson, editors, *Self-Organizing Architectures*, volume 6090 of *LNCS*, chapter 8, pages 152–176. Springer, July 2010. 1st International Workshop on Self-Organizing Architectures (SOAR 2009), Cambridge, UK, 14-17 September 2009, Revised Selected and Invited Papers.

[VCO09] Mirko Viroli, Matteo Casadei, and Andrea Omicini. A framework for modelling and implementing self-organising coordination. In Sung Y. Shin and Sascha Ossowski, editors, *SAC*, pages 1353–1360. ACM, 2009.

[VD14] Mirko Viroli and Ferruccio Damiani. A calculus of self-stabilising computational fields. In eva Kühn and Rosario Pugliese, editors, *16th Conference on Coordination Languages and Models (Coordination 2014)*, volume 8459 of *LNCS*, pages 163–178. Springer-Verlag, June 2014. Proceedings of the 16th Conference on Coordination Models and Languages (Coordination 2014), Berlin (Germany), 3-5 June.

[VDB13] Mirko Viroli, Ferruccio Damiani, and Jacob Beal. A calculus of computational fields. In Carlos Canal and Massimo Villari, editors, *Advances in Service-Oriented and Cloud Computing*, volume 393 of *Communications in Computer and Information Science*, pages 114–128. Springer Berlin Heidelberg, 2013.

[VHR+07] Mirko Viroli, Tom Holvoet, Alessandro Ricci, Kurt Schelfthout, and Franco Zambonelli. Infrastructures for the environment of multiagent systems. *Autonomous Agents and Multi-Agent Systems*, 14(1):49–60, July 2007.

[Vir11] Mirko Viroli. A self-organising approach to competition and composition of pervasive services. Technical Report SAPERE TR.WP2.2011.4, 2011.

[Vir13] Mirko Viroli. Engineering confluent computational fields: from functions to rewrite rules. In *Spatial Computing Workshop (SCW 2013)*, AAMAS 2013, Saint Paul, Minnesota, USA, May 2013.

[VMP+14] Mirko Viroli, Sara Montagna, Danilo Pianini, Graeme Stevenson, and Franco Zambonelli. A coordination model of pervasive service ecosystems. *Science of Computer Programming*, 2014. Submitted for publication.

[VO06] Mirko Viroli and Andrea Omicini. Coordination as a service. *Fundamenta Informaticae*, 73(4):507–534, 2006.

[VPB12] Mirko Viroli, Danilo Pianini, and Jacob Beal. Linda in space-time: an adaptive coordination model for mobile ad-hoc environments. In Marjan Sirjani, editor, *Coordination Languages and Models*, volume 7274 of *LNCS*, pages 212–229. Springer-Verlag, June 2012. Proceedings of the 14th Conference of Coordination Models and Languages (Coordination 2012), Stockholm (Sweden), 14-15 June.

[VPMS12] Mirko Viroli, Danilo Pianini, Sara Montagna, and Graeme Stevenson. Pervasive ecosystems: a coordination model based on semantic chemistry. In Sascha Ossowski, Paola Lecca, Chih-Cheng Hung, and Jiman Hong, editors, *27th Annual ACM Symposium on Applied Computing (SAC 2012)*, Riva del Garda, TN, Italy, 26–30 March 2012. ACM.

[VZ10] Mirko Viroli and Franco Zambonelli. A biochemical approach to adaptive service ecosystems. *Information Sciences*, 180(10):1876–1892, May 2010.

[VZSD12] Mirko Viroli, Franco Zambonelli, Graeme Stevenson, and Simon Dobson. From soa to pervasive service ecosystems: an approach based on semantic web technologies. In Javier Cubo and Guadalupe Ortiz, editors, *Adaptive Web Services for Modular and*

Reusable Software Development: Tactics and Solution, chapter 8, pages 207–237. IGI Global, 2012.

[Wal00] J. Waldo. *The Jini Specifications*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2nd edition, 2000.

[WBH06] Danny Weyns, Nelis Boucké, and Tom Holvoet. Gradient field-based task assignment in an agv transportation system. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, AAMAS '06, pages 842–849, New York, NY, USA, 2006. ACM.

[WBYI08] Justin Werfel, Yaneer Bar-Yam, and Donald Ingber. Bioinspired environmental coordination in spatial computing systems. In *1st International SASO Workshop on Spatial Computing*, pages 338–343, Washington, DC, USA, 2008. IEEE Computer Society.

[Weg97] Peter Wegner. Why interaction is more powerful than algorithms. *Commun. ACM*, 40(5):80–91, May 1997.

[WFR+12] Martin Wirz, Tobias Franke, Daniel Roggen, Eve Mitleton-Kelly, Paul Lukowicz, and Gerhard Troster. Inferring crowd conditions from pedestrians' location traces for real-time crowd monitoring during city-scale mass gatherings. *2012 IEEE 21st International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises*, 0:367–372, 2012.

[WFR+13] Martin Wirz, Tobias Franke, Daniel Roggen, Eve Mitleton-Kelly, Paul Lukowicz, and Gerhard Tröster. Probing crowd density through smartphones in city-scale mass gatherings. *EPJ Data Science*, pages 1–24, 2013.

[WL13] Jens Weppner and Paul Lukowicz. Bluetooth based collaborative crowd density estimation with mobile phones. In *IEEE International Conference on Pervasive Computing and Communications (PerCom)*, volume 18, page 22, 2013.

[WLS+02] Brian White, Jay Lepreau, Leigh Stoller, Robert Ricci, Shashi Guruprasad, Mac Newbold, Mike Hibler, Chad Barb, and Abhijeet Joglekar. An integrated experimental environment for distributed systems and networks. pages 255–270, Boston, MA, December 2002.

[WM04] Matt Welsh and Geoff Mainland. Programming sensor networks using abstract regions. In *Proceedings of the First USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI '04)*, March 2004.

[WMLF98] Peter Wyckoff, Stephen W. McLaughry, Tobin J. Lehman, and Daniel A. Ford. T Spaces. *IBM Systems Journal*, 37(3):454–474, 1998.

[WRG+07] Z. Wu, A. Ranabahu, K. Gomadam, A.P. Sheth, and J.A. Miller. Automatic composition of semantic web services using process and data mediation. In *Proc. of the 9th Intl. Conf. on Enterprise Information Systems*, pages 453–461, 2007.

[WSBC04] Kamin Whitehouse, Cory Sharp, Eric Brewer, and David Culler. Hood: a neighborhood abstraction for sensor networks. In *Proceedings of the 2nd international conference on Mobile systems, applications, and services*. ACM Press, 2004.

[Yam07] Daniel Yamins. *A Theory of Local-to-Global Algorithms for One-Dimensional Spatial Multi-Agent Systems*. PhD thesis, Harvard, Cambridge, MA, USA, December 2007.

[YDM12] Juan Ye, Simon Dobson, and Susan McKeever. Situation identification techniques in pervasive computing: A review. *Pervasive and Mobile Computing*, 8(1):36–66, 2012.

[Zam12] Franco Zambonelli. Toward sociotechnical urban superorganisms. *IEEE Computer*, 47(8), 2012.

[ZCF+11] Franco Zambonelli, Gabriella Castelli, Laura Ferrari, Marco Mamei, Alberto Rosi, Giovanna Di Marzo, Matteo Risoldi, Akla-Esso Tchao, Simon Dobson, Graeme Stevenson, Yuan Ye, Elena Nardini, Andrea Omicini, Sara Montagna, Mirko Viroli, Alois Ferscha, Sascha Maschek, and Bernhard Wally. Self-aware pervasive service ecosystems. *Procedia Computer Science*, 7:197–199, December 2011. Proceedings of the 2nd European Future Technologies Conference and Exhibition 2011 (FET 11).

[Zei76] Bernard P. Zeigler. *Theory of Modeling and Simulation*. John Wiley, 1976.

[Zei84] B. Zeigler. *Multifaceted modelling and Discrete Event Simulation*. Academic Press, 1984.

[ZLC+08] Yu Zheng, Quannan Li, Yukun Chen, Xing Xie, and Wei-Ying Ma. Understanding mobility based on gps data. In *In Proceedings of ACM conference on Ubiquitous Computing (UbiComp 2008)*, pages 312–321, Seoul, Korea, 2008. ACM Press.

[ZO04] Franco Zambonelli and Andrea Omicini. Challenges and research directions in agent-oriented software engineering. *Autonomous Agents and Multi-Agent Systems*, 9(3):253–283, 2004.

[ZOA+14] Franco Zambonelli, Andrea Omicini, Bernhard Anzengruber, Gabriella Castelli, Francesco L. DeAngelis, Giovanna Di Marzo Serugendo, Simon Dobson, Jose Luis Fernandez-Marquez, Alois Ferscha, Marco Mamei, Stefano Mariani, Ambra Molesini, Sara Montagna, Jussi Nieminen, Danilo Pianini, Matteo Risoldi, Alberto Rosi, Graeme Stevenson, Mirko Viroli, and Juan Ye. Developing pervasive multi-agent systems with nature-inspired coordination. *Pervasive and Mobile Computing*, 2014. Special Issue on "10 years of Pervasive Computing" in honor of Chatschik Bisdikian.

[ZRFS11] Kashif Zia, Andreas Riener, Alois Ferscha, and Alexei Sharpanskykh. Evacuation simulation based on cognitive decision making model in a socio-technical system. *Distributed Simulation and Real Time Applications, IEEE/ACM International Symposium on*, 0:98–107, 2011.

[ZV11] Franco Zambonelli and Mirko Viroli. A survey on nature-inspired metaphors for pervasive service ecosystems. *International Journal of Pervasive Computing and Communications*, 7(3):186–204, 2011.

List of Figures

List of Tables

Appendices

A
Proofs of Theorems

A.1 Discontinuity of Discrete-Valued Fields

Theorem 2. *Any field f with a discrete range and at least one event m with a neighbour m' such that $f(m) \neq f(m')$ cannot be continuous.*

Consider the neighbourhood $\mathscr{N}(m)$; by assumption, the set $f(\mathscr{N}(m))$ is finite and discrete, thus all of its subsets are open. If f is continuous, then both preimages $f^{-1}(f(m))$ and $f^{-1}(f(\mathscr{N}(m)) - f(m))$ are open. Their subsets intersecting with $\mathscr{N}(m)$ must also be open, and form a partition of $\mathscr{N}(m)$. Since $\mathscr{N}(m)$ is a connected subset of a manifold M, however, the complement of an open set is closed and not open. This can only be true if one of the two sets is empty, but since there is a neighbour m' mapping to a different value than m, we have a contradiction and f cannot be continuous.

A.2 Consistency of GPI-calculus

We prove consistency of GPI-calculus in three stages: First, we prove that any finite composition of eventually consistent and continuity preserving operators is also eventually consistent and continuity preserving. We then show that each operator in GPI-calculus is individually at least eventually consistent and continuity preserving. Finally, we show that all GPI-calculus programs are finite compositions of operators, which implies that all such programs are eventually consistent.

We begin by showing that eventual consistency and eventual continuity are preserved through composition:

Lemma 3. *Consider a set of operators that are eventually consistent and where, when there is a spatial section S_M such that environment e and inputs f_i are continuous on $e^{-1}(\mathbb{V} - \mathscr{B}) \cap T^+(S_M)$ and $f_i^{-1}(\mathbb{V} - \mathscr{B}) \cap T^+(S_M)$, then there must be a spatial section S_M such that the output f_o is continuous on $f_o^{-1}(\mathbb{V} - \mathscr{B}) \cap T^+(S_M)$ Any finite composition of instances of such operators has the same properties.*

Consider a well-defined program comprised of a finite sequence of operator instances, such that the outputs of some operator instances are inputs for others. Because the sequence is finite and the program is well-defined, it must be the case that the operators instances can be ordered, such that the ith operator instance depends only on the environment and the output fields of prior operators in the sequence.

For the *i*th operator instance in this sequence, it is possible to construct an equivalent program comprising only that operator instance, an environment e_i containing its inputs, and `sense` operator instances mapping the environment values to its inputs. As noted in Lemma 8, the `sense` operators trivially satisfy the consistency and continuity conditions, so the *i*th operator instance has its input conditions satisfied. Thus, by the assumption about operator properties, we also have that this miniature program is eventually consistent and that the continuity properties hold as well.

This output can then be added to the environment for the $i+1$ operator instance, ensuring that if the preconditions are satisfied for the *i*th operator instance, they will be satisfied for the $i+1$st operator instance as well. Thus, for any finite composition of operators, the eventual consistency and continuity property stated above holds for the composition if it holds for all of the individual operators[1].

We now move on to proving that each of the operators in GPI-calculus provides sufficient consistency and continuity, beginning with the trivial case of literals:

Lemma 4. *Literals* `l` *are consistent and have a continuous output.*

Trivially true, since the output field of a literal `l` is equal to `l` at every point in its domain.

Lemma 5. *Any built-in operators* `m` *is consistent and has output f_o continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})$ if its inputs f_i are approximable and continuous on $f_i^{-1}(\mathbb{V}-\mathscr{B})$.*

The output of any `m` operator at any event m is affected only by the values of its inputs at m. `m` is also continuous by definition, and (by the semantics of Field Calculus) requires all inputs to have the same domain. Thus, since the composition of continuous functions is continuous, when each input f_i is continuous on $f_i^{-1}(\mathbb{V}-\mathscr{B})$, it must be the case that the output f_o is continuous on $\bigcap_i f_i^{-1}(\mathbb{V}-\mathscr{B})$. The complementary space, $\bigcup_i f_i^{-1}(\mathscr{B})$ covers only points where at least one input has value $\mathscr{B}$, and thus by the definition of `m`, f_o maps all events in this space to $\mathscr{B}$.

Only consistency remains to be shown. All ε-approximation sequences must approximate f_o on subspace $f_o^{-1}(\mathbb{V}-\mathscr{B})$ because it is continuous on this space. The complementary space $\bigcup_i f_i^{-1}(\mathscr{B})$ must also always be approximated because it holds the constant value $\mathscr{B}$ on a finite union of subspaces.

Lemma 6. *Built-in operator* `mux` *is consistent and has output f_o continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})$ if its inputs f_i are approximable and continuous on $f_i^{-1}(\mathbb{V}-\mathscr{B})$.*

If `mux` is continuous on its first input $f_1^{-1}(\mathbb{V}-\mathscr{B})$, then the preimages of $f_1^{-1}(\mathit{true})$ and $f_1^{-1}(\mathit{false})$ must both be open. The output is then defined piecewise as:

$$f_o(m) = \begin{cases} f_2(m) & m \in field_1^{-1}(\mathit{true})) \\ f_3(m) & m \in field_1^{-1}(\mathit{false})) \\ \mathscr{B} & m \in field_1^{-1}(\mathscr{B})) \end{cases}$$

Since f_2 and f_3 are continuous on their non-$\mathscr{B}$ events, that must also hold for the output field on the open subspaces selected by the preimages of *true* and *false*. Likewise, any ε-approximation sequence

[1]Note that if the program were not guaranteed finite evaluation, then this result would not hold: the eventual consistency of the *i*th instance evaluated would still hold, but no *i* would be high enough to cover the entire sequence.

that approximates both f_2 and f_3 must also approximate the output field on those subspaces, because the values on each subspace are identical to one of the two inputs that is being approximated.

Finally, the space of events mapping to $\mathscr{B}$ in the output field must always be approximated as well because it holds the constant value $\mathscr{B}$ on a subspace constructed by finite intersection and union of subspaces.

Lemma 7. *Built-in operator < is consistent and has output f_o continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})$ if its inputs f_i are approximable and continuous on $f_i^{-1}(\mathbb{V}-\mathscr{B})$.*

The < operator only outputs three values, *true*, *false*, and $\mathscr{B}$, so this theorem will be satisfied if the preimages of both *true* and *false* are open.

We must consider two conditions: real number comparison and integer comparison. For each, we will describe the case only of events that f_o maps to *true*; the *false* case follows by symmetry.

For real number comparison, consider an event m mapping to *true*. This means that in the input fields, $f_2(m)-f_1(m)=\varepsilon>0$. Because the inputs are continuous at m, There must then be some δ, such that the open set of all events m' within δ of m also have difference $f_2(m')-f_1(m')>0$, meaning they also map to true. Any union of open sets is open, so $field_o^{-1}(true)$ must be open.

Integer comparison works the same way, except that the radius δ open sets around m are guaranteed to have inputs being equal to $f_1 m$ and $f_2 m$. Thus, it is also possible in the *false* case for the two input values to be equal and still produce an open preimage.

Lemma 8. *Built-in operator `sense` is consistent and has output f_o continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})$ if its environment is approximable and continuous on $e^{-1}(\mathbb{V}-\mathscr{B})$*

The environment is assumed to be a field mapping to a tuple of values at each point, and `sense` outputs a field created by indexing into said tuples by a literal integer. A field of tuples cannot be continuous unless each of its elements is also continuous. The output is then a copy of one of those elements, and so must be continuous on the same range.

Lemma 9. *Branch operator `if` is consistent and has output f_o continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})$ if its environment e and inputs f_i are approximable and continuous on $e^{-1}(\mathbb{V}-\mathscr{B})$ and $f_i^{-1}(\mathbb{V}-\mathscr{B})$ respectively.*

The `if` operator is identical to `mux` except that its branch expressions are evaluated with respect to the subspaces $f_1^{-1}(true)$ and $f_1^{-1}(false)$ respectively (i.e., the domain of the evaluation environment is reduced). Since these are open subspaces of the environment's domain M, the approximability and continuity properties of e are not affected by the domain reduction. Thus, if the branch expressions would have conformed with the preconditions for M they will also conform with the preconditions for the branch subspaces. The output field is then assembled piecewise identically to `mux`, and is approximable and continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})$ by the same reasoning.

Lemma 10. *Function call operator `f` is eventually consistent and has some spatial section S_M such that output f_o is continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})$ if its environment e and inputs f_i are approximable and continuous on $e^{-1}(\mathbb{V}-\mathscr{B})$ and $f_i^{-1}(\mathbb{V}-\mathscr{B})$ respectively.*

Consider a function definition `f`; either `f` contains other function calls or it does not. If it does not, then evaluating `f` is equivalent to evaluating a composition of operators satisfying Lemma 3 (adjoining the

function arguments to the environment and substituting `sense` functions for variable references). Thus the desired consistency and continuity properties hold.

If `f` does contain function calls, then the properties hold if they hold for all of the function calls within `f` (meaning that once again Lemma 3 can apply). Since GPI-calculus does not allow recursion, it must be the case that these dependencies between function definitions can be arranged in a finite directed acyclic graph. The nodes of such a graph may then be ordered, such that each subsequent node only depends on nodes before it in the order. Since the set is finite, there must be at least one node that has no dependencies and thus forms a base case for induction showing that the properties hold for all function calls.

Lemma 11. *Operator `GPI` is eventually consistent and has output f_o continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})$ in the future of some spatial section S_M if its environment e and inputs f_i are approximable and continuous in the future of some spatial section S_M on $e^{-1}(\mathbb{V}-\mathscr{B})$ and $f_i^{-1}(\mathbb{V}-\mathscr{B})$ respectively.*

Following the semantics of Field Calculus to interpret the `GPI` algorithm given in Section 3 of the main text, the first element of the tuple computed in the `rep` statement implements a computation of distance via the triangle inequality (`nbr-range` is a metric, and a metric multiplied by a continuous positive scalar function is still a metric).

Thus, if there is a spatial section S_M such that the values of all of the inputs do not change at any device on $T^+(S_M)$, then since we assume that manifolds have finite diameter, it must be the case that all the distance estimates (first values of the `distance-integral` tuple) eventually converge to a continuous field of distance estimates. The values of the integral are co-computed with the values of the distance estimates, so they too will stop changing once the inputs have stopped changing. Thus it must be possible to choose a spatial section of S_M on which values of `distance-integral` do not change at any device.

Note also that due to the use of `min-hood`' in computing the triangle inequality, any point in the field of distances with more than one shortest path leasding to it will be replaced by a $\mathscr{B}$. This eliminates only a set of measure zero: because the distance function is continuous, its gradient cannot be discontinuous on a space of more than measure zero. The set eliminated is, however, precisely the set of points for which the gradient of the distance field is not continuous.

This is important because this is also the set of points on which the integral calculated in the second element of the `distance-integral` tuple might not be continuous. Consider, for example, the case shown in Figure A.1, in which a hole causes there to be two very different shortest paths to a point. The integrals computed along such paths may be very different indeed, necessarily creating a discontinuity in the value of the integral and hence the output of `GPI`. This type of problem can also come from other sources besides topological complexity: similar patterns (and failures) can be caused by `if` statements, distortions in the distance measure, or the shape of the source—anything that can cause a discontinuity in the gradient.

Since `min-hood`' ensures that such points are $\mathscr{B}$, however, they are eliminated from the region on which we must establish continuity. If we instead consider any m with precisely one shortest path to the `source` region, then because both integrand and the gradient of the computed distance field are continuous on this path, it must be the case that for any given ε, there must be a δ such that the open set of events within distance δ have integral values that are less than δ different, and thus the field of integrals output by `GPI` is continuous on $f_o^{-1}(\mathbb{V}-\mathscr{B})\cap T^+(S_M)$. Because it is continuous on this space, it must also be approximable; the complementary space of $\mathscr{B}$ values must also be approximable, as it is a union

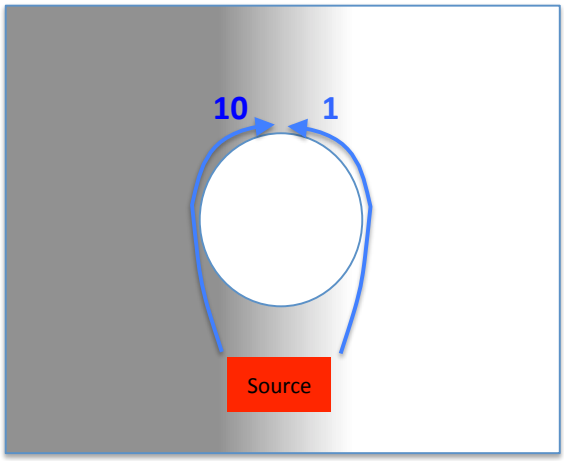

Figure A.1: The field of path integral values created by `GPI` can be discontinuous at points reached by two different shortest paths, as in this example where the left path around the hole goes through a region where the integrand is much higher than on the right path, resulting in discontinuity between a value of 10 from the left and 1 from the right.

of the $\mathscr{B}$ values of the inputs (which must be approximable) plus the measure zero set of $\mathscr{B}$ events added from computation of the distance function.

Theorem 12. *GPI-calculus programs are eventually consistent for all environments e that are continuous on $e^{-1}(\mathbb{V}-\mathscr{B})$.*

By Lemmas 4 through 11, we know that every operator in GPI-calculus has the property that if its inputs f_i and environment e are continuous on $f_i^{-1}(\mathbb{V}-\mathscr{B})$ and $e^{-1}(\mathbb{V}-\mathscr{B})$ respectively, then it is eventually consistent (all consistent programs are of course also eventually consistent). We further know that, when there is a spatial section S_M such that inputs are approximated on $T^+(S_M)$, there must also be some spatial section S_M such that the output is approximated on $T^+(S_M)$ and continuous on $T^+(S_M)\cap f_o^{-1}(\mathbb{V}-\mathscr{B})$ (note that operators continuous $f_o^{-1}(\mathbb{V}-\mathscr{B})$ are also continuous on open subsets thereof, e.g., $T^+(S_M)$).

We now consider a program as a finite sequence of operator instances. Because recursion is prohibited, it must be the case that the number of operators evaluated in the evaluation of a GPI-calculus program must be finite for any given ε-approximation. Since `if` is the only method of branching evaluations, it must thus be the case that in any approximation sequence there is some i, after which no $\varepsilon_{j>i}$-approximation evaluates an operator instance that is not also evaluated in some prior ε-approximation, considering any instance in which an operator instance is not evaluated to be an evaluation with null domain.

Given the assumed base case of an environment continuous on $e^{-1}(\mathbb{V}-\mathscr{B})$, we thus satisfy the conditions of Lemma 3 and have that the program must be eventually consistent.