

A Bio-chemical Approach to Awareness in Pervasive Systems

Graeme Stevenson, Juan Ye
and Simon Dobson
School of Computer Science
University of St Andrews, UK
graeme.stevenson@st-andrews.ac.uk

Gabriella Castelli, Alberto Rosi
and Franco Zambonelli
Dept. of Science and Engineering Methods
University of Modena and Reggio Emilia, Italy
firstname.lastname@unimore.it

ABSTRACT

The emerging proliferation of devices with sensing and networking technologies offers opportunities for delivering pervasive services through interactions between spatially local resources. Here, applications are challenged to become aware of their surroundings: to discover, filter and reason on information relevant to their goals. Without centralised services to control information flow, decentralised mechanisms must partition these responsibilities across the environment. This paper explores a bio-chemically inspired approach to realising awareness in an open, dynamic setting, building towards complex self-organising awareness algorithms for data collection, reasoning, and querying. Through simulation we provide a preliminary evaluation of the proposed approach.

Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: Search process, Information filtering; H.3.4 [Systems and Software]: Distributed systems, Performance evaluation

General Terms

Design, Experimentation, Measurement, Performance

Keywords

Pervasive systems, Situation awareness, Bio-chemical

1. INTRODUCTION

A recent survey [11] predicts that by 2016 over one billion people will own mobile devices, each embedded with powerful sensing, computing, and networking capabilities. This creates an exciting opportunity for harnessing very large volumes of sensed information to support the delivery of context-aware services to human users. From digital signs that adapt to their audiences' preferences [12] to the inference of friendship networks from mobile phone data [5], glimpses of how such advances may soon be fully exploited are numerous.

Contrasting with traditional service-oriented systems where physical situation and spatiality are not primary abstractions, we envision a point in the near future where devices will opportunistically and spontaneously connect with each other to form dense infrastructures, *pervasive ecosystems*. Here, resources are exploited to offer spatially situated services that self-organise to their context of use, enhancing service delivery with myriad sensing modalities offered by the surrounding physical and virtual world [17].

The path towards realising this vision is paved with substantial technical challenges to support services in becoming aware of their surroundings: to seek, discover and filter information relevant to their goals, and to aggregate, interpret and reason over such information to identify *situations* – semantically meaningful classifications of states – of importance upon which they will adapt their behaviour. These scenarios cannot assume the existence of centralised services to which all the data in the network is fed and from which classifications of situations of interest subsequently flow. Consequently, pervasive ecosystems require mechanisms to partition these responsibilities across the environment's participants, with awareness arising from the interactions between many individuals, each exposing aspects of their own state and perceiving the states of others to further increase their awareness of their surroundings.

This paper approaches awareness in an open, dynamic, and mobile environment following the natural inspiration, where we exploit a bio-chemical model to support spontaneous and self-organising interactions between entities. Participant entities, which can be people, their devices, software services, and physical and virtual resources, are modelled heterogeneously by continuously updated textual *annotations* that describe aspects of their goals and beliefs and state their wish to make public.

The environment embeds the basic “laws of nature” that govern the ecosystem, manipulating annotations like reagents in a chemical reaction to regulate the interactions between annotations, relate them dynamically based on content, synthesise new information from old, and manage the diffusion of annotations across the environment. Entities only interact with the ecology through manipulating and observing the structure and content of their annotations.

This vision is being developed in the SAPERE project [6], wherein we are realising a lightweight middleware to enable

SENSEMINI'13 November 14 2013, Roma, Italy.

Copyright is held by the owner/author(s).

Publication rights licensed to ACM.

ACM 978-1-4503-2430-4/13/11...\$15.00.

<http://dx.doi.org/10.1145/2536714.2536721>

PCs, notebooks, modern tablets and smartphones to be part of the environment, all embedding the above functionalities.

The novelty of the work described here resides in: (i) exploiting locally-scoped interactions inspired by the bio-chemical metaphor to elaborate a spectrum of awareness mechanisms accounting for increasing levels of complexity (Section 2); and (ii) developing self-organising awareness algorithms by combining spectrum mechanisms, providing a basis for situation recognition (Section 3). Simulations preliminarily evaluate the proposed algorithms as a means of supporting scalable decentralised awareness (Section 4), and we outline related work (Section 5) before concluding (Section 6).

2. A SPECTRUM OF AWARENESS

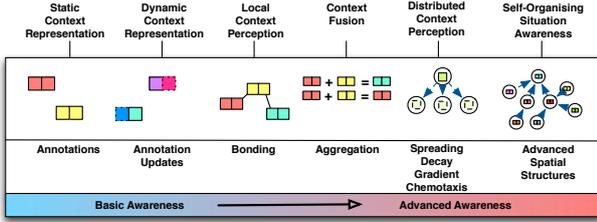


Figure 1: The awareness spectrum.

Bio-chemical self-* mechanisms have been successfully applied in different domains, achieving results that go beyond traditional approaches. In [7] the authors illustrate a catalogue of bio-chemical mechanisms that are used in this paper and form the building blocks for engineering behaviours in a self-organising system. Through their application, information embedded in annotations can be made available to and observed by other participants, perceived over regions of space, and fused with other information.

By composing these bio-chemical mechanisms, more complex interactions and communication strategies emerge, enabling increasingly complex forms of awareness to be realised within the ecosystem. Information can be sought by and delivered to a requestor, and complex spatial data structures can be constructed to support the self-aggregation and self-composition of data [8] for the purpose of higher-level, unsupervised reasoning. We view the types of awareness attainable as a spectrum, illustrated in Figure 1, with basic data representation at one end and high-level awareness achieved through the bio-chemical collaborations of ecosystem participants at the other. From left to right we mark six points along this continuum and briefly overview the awareness capabilities available at each.

Static Context Representation The representation of static state is a basic need of a context aware system, encompassing any information unlikely to change state, such as a building map or historic information. Such annotation can be made available through a tuple A , that at a location l , with content x , is a collection of key-value pairs:

$$A^l[x], x = (k_1 = v_1, k_2 = v_2, \dots, k_n = v_n) \quad (\text{STATIC})$$

Dynamic Context Representation Annotations provide mutable fields capable of expressing and seamlessly updating

information about internal states, goals, desires and beliefs, thus the value of k_j at time t , v_j , is represented as:

$$A^l[x], x = (\dots, k_j = v_j(t), \dots) \quad (\text{DYNAMIC})$$

Local Context Perception The bio-chemical mechanism of *bonding*, a virtual chemical-like bond between two annotations, provides participants with the means to inspect the static and dynamic information embedded in the annotations of co-located individuals. Informally, a bond (\oplus) between two co-located annotations (denoted by the shared superscript l) A and B , is created (\longrightarrow) when the former expresses a request x , that is satisfied by y in the later:

$$A^l[x] + B^l[y] \longrightarrow A \oplus B, \quad (\text{BOND})$$

Context Fusion Context awareness is aided by fusing information to attain higher level insights. It is supported through the mechanism of *aggregation*, which works by replacing two annotations (and through iterative application, many) with a single summarising annotation. It is a primary means of carrying out selection and fusion, and has the effect of reducing the overall information in the environment.

The general form of the mechanism is illustrated below, where two annotations A and B with content x and y , are removed from the environment and replaced by an annotation C . The content description identifies both shared and individual properties that control which annotations the reaction is applied to, and which parts of the content are affected. In its trivial form, C represents an unmodified version of A or B , for example, where the most recent version of an annotation is retained in preference to an older version.

$$A^l[x] + B^l[y] \longrightarrow C^l \quad (\text{AGG})$$

Distributed Context Perception Next, we use *spreading* to support ecosystem participants in perceiving annotations in their non-immediate proximity. Spreading works by replicating an annotation across neighbouring devices, bounded by a metric such as hops, or physical distance. We represent this mechanism below, with application of the mechanism resulting in the creation of a copy of annotation A at location l' , without modification to the original.

$$A^l[x] \longrightarrow A^l + A^{l'}, \text{ until boundary} \quad (\text{SPREAD})$$

The spreading mechanism allows participants to project their state across a region of the environment by setting properties in their annotation to trigger the reaction. Participants may bond with spread annotations as they would any local annotation to perceive their content and act accordingly. To manage the lifecycle of distributed information, all spread data is associated with a counter c , with the *decay* mechanism decrementing this counter over time, leading to the annotation's eventual removal when the counter reaches zero:

$$A^l[c = i] \longrightarrow A^l[c = i - 1] \quad (\text{DECAY}_1)$$

$$A^l[c = 0] \longrightarrow \quad (\text{DECAY}_2)$$

By composing the four basic mechanisms that support awareness it becomes possible to support the structured search

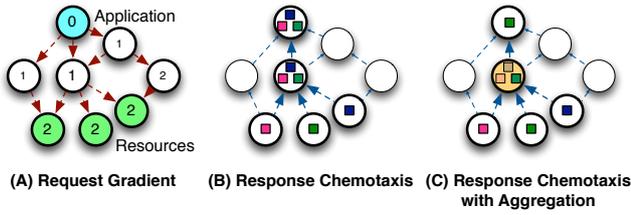


Figure 2: Context acquisition interactions.

for and directed propagation of annotations. Two particular combinations – *gradient* and *chemotaxis* – provide this functionality. The gradient mechanism provides an effective way to acquire local awareness by the periodic update and aggregation of annotations reflecting pointers towards a source’s physical position spread across the environment, while chemotaxis uses gradients to enable information routing via the shortest accessible path. These mechanisms provide the final self-organising foundations for realising complex awareness mechanisms within our bio-chemically inspired pervasive ecosystem model.

Distributed context awareness is useful to seek or project the presence of resources across a region of space—informing interested parties of state, advertising services, and serving as the basis for further interactions. To make the discussion more concrete, consider a congestion monitoring application at a city festival. Classifier nodes request information about the current position of entities over a region of the city through a set of gradient embedded requests. On receipt of the gradient annotation by a car, bus, or smartphone, an annotation describing the entity’s positions and trajectory is routed towards the classifier node using chemotaxis, where it can be further processed. For example, if the congestion level in a particular region is found to exceed a threshold, actions, such as requesting that people move away from an area, can be taken. Figure 2 illustrates the key steps: (A) an application constructs an annotation with an embedded request and sets the necessary properties to trigger construction of a gradient; spreading and aggregation patterns establish a gradient data structure in which each annotation has a pointer to the node in which it was created, which is also the node indicating the direction to the source following the shortest path; (B) where a request annotation is co-located with another annotation, matching is applied to resolve their compatibility; a successful match generates a reply annotation pointing towards the request; (C) aggregation may be applied to reduce transmission costs or selectively filter replies *en route* to a common destination [16].

Self-Organising Situation Awareness We believe that the spectrum of awareness mechanisms are general enough to support a wide range of context awareness applications. The final point on the spectrum, to the far right-hand-side, represents any general, but complex awareness behaviour that can be thought of as providing advanced self-organisation; that is, behaviour that exhibits little-, or, ideally, no-dependency to the presence of specific ecosystem participants, and can adapt to the addition and removal of role-playing participants appropriately. Such algorithms are the topic of the next section.

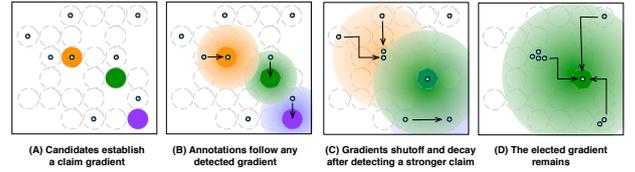


Figure 3: The election process that forms part of self-organising data collection.

3. SELF-ORGANISING AWARENESS

We now describe one possible approach to (A) self organising data collection, (B) data segmentation, and (C) semantic clustering and querying, that uses awareness spectrum techniques introduced above to achieve self-organising situation awareness. In the following, by describing how they may be realised in a pervasive ecosystem, we seek to demonstrate not only the flexible and extensible nature of the types of awareness that may be supported, but also the suitability of the bio-chemical substrate for approximating what can be thought of “traditional”, non bio-chemical, algorithms.

3.1 Self-Organising Data Collection

To gain awareness through reasoning over multiple sources, information be drawn towards a common point in the network. Often, it can be desirable to perform this pooling or collection of data in a self-organising way.

Consider again the example of a traffic and pedestrian congestion monitoring in a city. Self organising data-collection can be used to abstract and summarise the data generated across a physical region of space from both fixed sensors deployed in the environment (e.g., traffic or CCTV cameras) and as well as from mobile devices (pedestrian position, density, and speed) where no fixed infrastructure exists. Complicating this process is the notion that the set of devices active in a region of space change over time; therefore it follows that the selection of the devices within which to perform such computation should be dynamic, self organising, and self healing in response to host mobility, failure, or the introduction of new devices into the region.

Self-organising data collection may be achieved, combining local and distributed context perception to elect a host responsible for the process, as illustrated in Figure 3:

1. Employ local context perception to gather knowledge of the host credibility for supporting data collection;
2. Express the local host’s credentials as an annotation. For example, a host’s id, connectivity, mobility, computational resources, or its physical position, are well-established heuristics for this purpose [13];
3. Employ distributed context perception to carry out the election of a host:
 - 3.i Broadcast devices’ claims via gradients;
 - 3.ii Switch off the claim-gradient if a better claim from a competing host is perceived locally, allowing the local claim to decay;
 - 3.iii Have annotations ascend the gradient structure; which will lead to the elected host on completion.

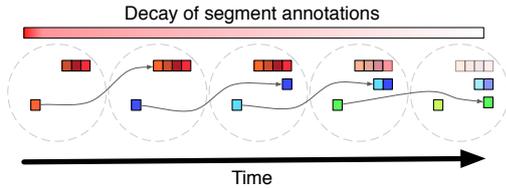


Figure 4: Segmentation illustrated, with colour and opacity analogues for semantic closeness and decay.

3.2 Data Segmentation and Correlation

After electing a host to collect data, pattern segmentation aims to divide sensor information (expressed as annotations) into semantically or structurally meaningful sequences. Following the example of a monitoring system, as data is collected over time, online data segmentation can be used to identify congestion patterns for different areas of the city across different days, times and regular events (e.g., shopping hours, sporting events and city festivals).

In most situation recognition approaches, segmenting sensor traces offline is a prerequisite step; for example, centrally collecting and then partitioning sensor traces in one-minute windows. Here we consider a novel technique that partitions sensor traces based on the semantic similarity between two adjacent sensor events; that is, if the similarity is below a given threshold, then the sensor events are grouped; otherwise, they are partitioned into two sequences. This method enables real-time, self-organising segmentation, based on the assumption that semantically similar events are meaningfully related and can be used as basis for situation recognition.

The segmentation process is a combination of distributed context perception and context fusion, as illustrated in Figure 4 and described below:

1. An agent starts the process by injecting an active segment annotation (simply, a property *active = true*);
2. Aggregation is triggered by the co-location of an active segment and any new sensor annotation arriving in the space:
 - 2.i If the sensor annotation is semantically similar to the segment annotation and occurs within some temporal distance of the previous event, the segment annotation is modified to incorporate the new information;
 - 2.ii Otherwise the segment annotation is tagged as inactive (*active = false*) and a new active segment annotation embedding the sensor annotation data is created to host a new sequence.
3. Decay is applied to segment annotations to remove old sequences over time.

3.3 Semantic Clustering and Querying

Having executed the segmentation process over a period of time, it is possible to abstract this information into a set of concise, meaningful, and frequent patterns that may indicate a state or situation of interest. In the context of the monitoring system example, semantic clustering can then be used to identify significant patterns, with live data then classified

against these patterns and queried to reveal whether current congestion levels match an expected pattern, or whether some unexpected event (e.g., a crash or a fire) is occurring may require investigation.

The process is performed by an agent binding to the segment annotations, whose key steps are illustrated as follows: (diagrammatically in Figure 5):

1. Bind to the set of segment annotations in a space;
2. Apply a k -means clustering algorithm:
 - 3.i Choose k patterns as the initial centroids;
 - 3.ii Assign each of the remaining patterns to one of these k groups if its distance to its centroid is closest using a semantic string alignment algorithm;
 - 3.iii Choose the pattern with the highest occurrence ratio in each group as new centroids;
 - 3.iv Repeat steps (3.ii) and (3.iii) until convergence.

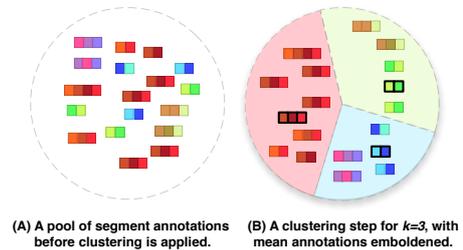


Figure 5: Illustrating clustering with $k = 3$ on a set of semantically segmented sensor trace annotations.

Performed periodically, with frequency dependent on how the underlying data evolves, it is a straightforward process to then match the live sensor data sequences (produced by the agent responsible for segmentation) to the discovered patterns by choosing the pattern most similar to the live sensor trace. We note that clustering is only one means of extracting patterns from data, with other machine learning techniques also applicable here.

Live querying of the identified situations is achieved similarly to regular resource discovery, either through first labelling the discovered semantically meaningful patterns, or by applications providing “template” descriptions of situations of interest, which can be matched to the discovered patterns similarly to the live data.

4. EVALUATION

Here, we present a preliminary evaluation of the performance of *context acquisition* and *self-organising data collection* as a starting point towards illustrating the potential scalability of the approach to awareness we propose. We report the recognition accuracy of a similar semantic technique applied to centralised, smart-home environments in [20].

This evaluation is conducted using The *ONE* [9], an opportunistic network environment simulator. An open environment of size $1km^2$ is constructed, with a set of homogeneous mobile hosts deployed. Each host runs the SAPERE middleware [1] and travels by means of a random walk, $0 - 4km/h$

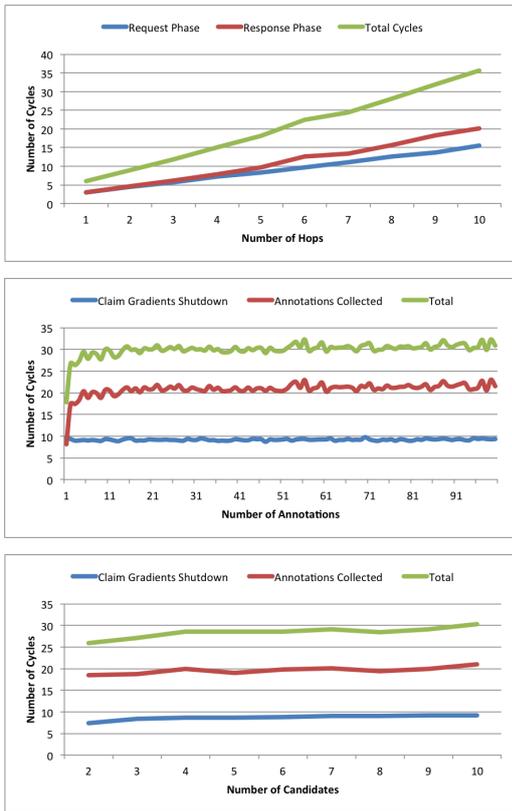


Figure 6: The linear scalability of context acquisition phases (top). Self-organising data collection with variable number of annotations (middle) and candidates (bottom).

and wireless communication range of 95m. To simulate a dense network infrastructure, as a city centre, we extrapolated a “necessary” network density of 400 devices per km^2 , assuring that a path between two randomly chosen hosts in the space exists with 99% probability. Analysing the performance on these algorithms in the presence of segmentation and high churn is left for future investigation.

4.1 Performance of the Awareness Algorithms

In this section we measure the performances associated with the execution of *context acquisition*, and *self-organising data collection* at the base of other described techniques. Each experiment is executed 1000 times over randomly generated topologies, with the results averaged.

Figure 6 (top) charts the performance of the request and response phases of the context acquisition process against the average number of cycles of the middleware engine (an execution each of the bond, aggregation, spread, and decay mechanisms against the local annotation set), grouped by the initial hop distance between requestor and responder. Both mechanisms perform similarly, the response phase performance slightly reduces at larger hop distances due to disruption to the chemotaxis trail, causing a delay until the gradient is next refreshed. The key result is that the total time for context acquisition can be seen to scale linearly with the hop-distance between the requestor and the responder;

implying scalability as the device density further increases.

Next, we investigate the performance of the self-organising data collection algorithm described in Section 3. We randomly position candidate nodes and annotations to be collected within a region extending 5-hops from the device that will be eventually elected and investigate the time for the election to complete (all losing gradients switch off), and the time for all the annotations to be collected by the elected device after competing candidate gradients have decayed. Active candidates refresh their gradients on each cycle, remaining active for 10 additional cycles. Figure 6 (middle and bottom) charts the performance of the self-organising data collection algorithm where both the number of annotations and the number of candidate devices are varied. In the first case the number of candidates is fixed at 5, while in the second case the number of annotations is fixed at 10.

The results show that the performance of self-organising data collection is invariant to both these factors. Referring to the resource collection performance for the same hop distance (see Figure 6), we see that main factor affecting gradient switch-off performance is the hop-distance between the elected host and the most distant candidate, while annotation collection incurs a delay corresponding to the number of cycles for the candidate gradients to decay after being switched off. This result, along with the results of evaluating the other awareness algorithms in this section imply good scalability characteristics for our approaches to realising self-organising distributed awareness in pervasive ecosystems.

5. RELATED WORK

Many approaches to awareness, such as iQL [3] and Construct [15], derive directly from The Context Toolkit: a highly influential framework supporting context awareness in small scale environments [4]. In these frameworks, applications use a well known discovery service to select and combine resources (e.g., sensors or aggregators) in a pipe-and-filter style according to semantic descriptions of their required inputs and outputs. Our approach to awareness differs from these approaches in four main respects: (i) it is designed to operate in a mobile environment; (ii) it treats spatial awareness and locality of interaction as being intrinsic; (iii) it does not rely on well known discovery services; and, (iv) it can be thought of as providing “intelligent” communication—delivering self-organising routing and information aggregation as part of the communication substrate, rather than through the use of “dumb” pipes.

Agent-based Context Aware Infrastructure (ACAI) [10], Co-BrA [2], and Gaia [14] are three approaches that allow context to be collected, processed, reasoned on, and disseminated to applications within smart spaces. These approaches are bound with physical locations - homes, offices, meeting rooms – and in the cases of ACAI and CoBrA, the interaction model follows a blackboard-style approach, which has limited potential to scale. Our approach considers the environment as an active participant, supporting the self-organisation of data acquisition, aggregation, and dissemination in a high-level way.

Finally, there is emerging work on awareness arising from the interactions between co-located devices. Weppner and

Lukowicz describe a technique for estimating crowd density through collaborative Bluetooth scanning [18], while Wirz et al. approach the detection of collective behaviour patterns of people from body worn sensors [19]. We believe these and similar works are expressible as self-organising situation awareness algorithms, as advocated by this paper.

6. CONCLUSION

This paper explored a bio-chemically inspired approach to realising application awareness in the setting of an open, dynamic, mobile and unpredictable operating environment. We described how incrementally increased awareness levels can be achieved through a small set of locally-scoped mechanisms, leading to complex self-organising awareness algorithms for data collection, reasoning, and querying. Finally, we presented a preliminary evaluation of our proposed approach through simulation.

Our future research will proceed in three main directions to: (i) investigate the role of awareness in self-managing the tuning of the bio-chemical mechanisms; (ii) investigate where lightweight bio-inspired mechanisms can play additional roles in the situation recognition process; and (iii) evaluate the mechanisms using real-world urban datasets to better evaluate real-time responsiveness, robustness, and adaptation to dynamic changes of context.

Acknowledgments

This work is supported by the Self-Aware Pervasive Service Ecosystems project (EU FP7-FET, Contract No. 256873).

7. REFERENCES

- [1] G. Castelli, M. Mamei, A. Rosi, and F. Zambonelli. Pervasive middleware goes social: the SAPERE approach. In *IEEE International Workshop on Socio-Aware Networked Computing Systems*, Ann Arbor, USA, 2011.
- [2] H. Chen, T. Finin, and A. Joshi. An intelligent broker for context-aware systems. In *Adjunct Proceedings of Ubicomp*, pages 183–184, October 2003.
- [3] N. H. Cohen, H. Lei, P. Castro, J. S. Davis II, and A. Purakayastha. Composing pervasive data using iQL. In *Fourth IEEE Workshop on Mobile Computing Systems and Applications*, pages 94–100, Callicoon, NY, USA, June 2002.
- [4] A. K. Dey. *Providing Architectural Support for Building Context-Aware Applications*. PhD thesis, Department of Computer Science, Georgia Institute of Technology, 2000.
- [5] N. Eagle, A. S. Pentland, D. Lazer, and S. Hanson. Inferring friendship network structure by using mobile phone data. volume 106, pages 15274–15278, 2009.
- [6] F. Zambonelli et al. Self-aware Pervasive Service Ecosystems. *Procedia Computer Science*, 7:197–199, 2011.
- [7] J. L. Fernandez-Marquez, G. Di Marzo Serugendo, S. Montagna, M. Viroli, and J. L. Arcos. Description and composition of bio-inspired design patterns: a complete overview. *Natural Computing*, pages 1–25, 2012.
- [8] J. L. Fernandez-Marquez, G. Stevenson, A. E. Tchao, J. Ye, G. Di Marzo Serugendo, and S. Dobson. Analysis of new gradient based aggregation algorithms for data-propagation in distributed networks. In *1st International Workshop on Adaptive Service Ecosystems: Natural and Socially Inspired Solutions*, Lyon, France, September 2012.
- [9] A. Keränen, J. Ott, and T. Kärkkäinen. The ONE Simulator for DTN Protocol Evaluation. In *2nd International Conference on Simulation Tools and Techniques*, Rome, Italy, 2009.
- [10] M. Khedr and A. Karmouch. ACAI: agent-based context-aware infrastructure for spontaneous applications. *Journal of Network and Computer Applications*, 28(1):19–44, 2005.
- [11] J. C. McCarthy, C. Mines, P. Matzke, and Y. Darashkevich. Mobile app internet recasts the software and services landscape. *Forrester Research Consumer PC And Tablet Forecast, 2011 to 2016*, 2011.
- [12] J. Müller, J. Exeler, M. Buzeck, and A. Krüger. Reflectivesigns: Digital signs that adapt to audience attention. In *7th International Conference on Pervasive Computing*, pages 17–24, May 2009.
- [13] Ramesh and Somasundaram. A Comparative Study Of Clusterhead Selection Algorithms In Wireless Sensor Networks. *International Journal of Computer Science & Engineering Survey*, 2(4):153–164, 2011.
- [14] A. Ranganathan and R. H. Campbell. A middleware for context-aware agents in ubiquitous computing environments. In *Proceedings of the 2003 International Conference on Middleware*, pages 143–161, Rio de Janeiro, Brazil, June 2003.
- [15] G. Stevenson, P. Nixon, and S. Dobson. Towards a reliable wide-area infrastructure for context-based self-management of communications. In *Second international IFIP conference on Autonomic Communication*, number 3854, pages 115–128, October 2005.
- [16] G. Stevenson, D. Pianini, S. Montagna, M. Viroli, J. Ye, and S. Dobson. Combining self-organisation, context-awareness and semantic reasoning: the case of resource discovery in opportunistic networks. In *28th Annual ACM Symposium on Applied Computing*, Coimbra, Portugal, March 2013.
- [17] M. Viroli, F. Zambonelli, G. Stevenson, and S. Dobson. From SOA to pervasive service ecosystems: An approach based on Semantic Web technologies. In *Adaptive Web Services for Modular and Reusable Software Development: Tactics and Solution*, chapter 8, pages 207–237. IGI Global, 2012.
- [18] J. Weppner and P. Lukowicz. Collaborative Crowd Density Estimation with Mobile Phones. In *Second International Workshop on Sensing Applications on Mobile Phones*, pages 1–5, Seattle, USA, 2011.
- [19] M. Wirz, D. Roggen, and G. Troster. Decentralized Detection of Group Formations from Wearable Acceleration Sensors. In *International Conference on Computational Science and Engineering*, pages 952–959, Vancouver, August 2009.
- [20] J. Ye and G. Stevenson. Dynamically detecting interleaved and concurrent activities. In *Proceedings of the International Conference on Ambient Intelligence*, 2013. To appear.