Knowledge Networks: the Nervous System of an Autonomic Communication Infrastructure

Maurice Mulvenna[1], Franco Zambonelli[2,] Kevin Curran[1], Chris Nugent[1]

[1]School of Computing and Mathematics, University of Ulster, Newtownabbey, BT37 0QB, UK
{md.mulvenna, kj.curran, cd.nugent }@ulster.ac.uk
[2] DISMI - Università di Modena e Reggio Emilia, Italy
franco.zambonelli@unimore.it

For future network scenarios to exhibit autonomic behaviour, both networks and application components and services need to be aware of their computational and environmental context, and must tune their activities accordingly. In this position paper, we propose an abstract architecture for knowledge networks that addresses the key issues of how both physical contextual knowledge and social knowledge from the users of communication networks can be used to form a knowledge space in support of autonomic agents dealing with network elements and applications. We discuss that the availability of raw contextual data is not enough to achieve meaningful autonomic behaviours. Rather, contextual information should be properly organised into 'networks of knowledge', to be exploited by both network and application components as the basic 'nervous system' in which situational stimuli reify into digital knowledge, and by means of which components can properly orchestrate their activities in a globally meaningful way. Here we firstly discuss the fundamental role of knowledge networks, and try to sketch what actual form and position such knowledge networks could assume. Then, we analyse some simple scenarios of use, showing how it is possible for the components of an autonomic communication system to build such knowledge networks autonomously; and, at the same time, to exploit them for orchestrating their activities in a type of stigmergy-based knowledge-rich system. Eventually, we sketch a rough research agenda and discuss the relations with other research areas.

1. Introduction

We envision that future networks will be able to provide composite, highly distributed, pervasive services in a situated and fully autonomic way. In other words, they will be made up of components capable of [KepC03, Zam05] understanding the general context – physical, technological, social, user-specific and request-specific – in which they operate; and spontaneously aggregating with each other and orchestrating their activities accordingly to that context, so as to support a range of activities and services activities that are simply not possible or impractical now, with the important addition of requesting no configuration efforts from users.

In particular, we expect services to be able to:

(i) Improve our interactions with the physical world by providing us with any needed information about our surrounding physical environment and exploiting such information to adapt/enrich their behaviour on such basis (e.g., consider adapting the behaviour of a tourist service network on the basis of the location from which the service is invoked and of the current weather and traffic conditions) [Est02];

(ii) Get the best of the network infrastructure and resources upon which they operate, being able to ensure sufficient quality of service adaptively and independently of the actual network characteristics (e.g., independently of the fact that we require them from a Wi-Fi PDA, from a GPRS phone, or from whatever connectivity and connected devices will be available at that time) [MikM04];

(iii) Facilitate our social interactions, by properly reflecting and exploiting the social context in which we are currently employing a service, e.g. for mere entertainment, or socialisation, or in the context of business activities. Such social possibilities could be particularly appreciated in an increasingly open and multicultural environment such as the EU [ChoP03, Pen05].

A central challenge for the above vision to become real is the promotion of suitable solutions for enabling the components of an autonomic communication infrastructure (whether network-level or application-level components) to become situation-aware. Assuming that mechanisms exists to produce all necessary "situational" knowledge (e.g., sensors and monitoring mechanism [Est02, Gel02], user and social profilers [Pen05], etc.), for components to exploit the knowledge properly it is necessary that all the available knowledge (which can be in a dramatic amount, can be distributed, decentralized, and can come from a multitude of sources) is organised for utilisation.

Organising all available situational information implies that any relations between information is properly represented and correlated (according to well-defined ontological constructs), so as to facilitate their retrieval and their understanding. To promote accessibility, it is necessary that information produced locally at one place is properly diffused in the network whenever this may be of a more global relevance. Also, it may be important that such information can be exploited for mediated (i.e., stigmergic) interactions among the components of the infrastructure, so as to promote both robust self-organising behaviours [DiM04] and fruitful cross-layer interactions.

These needs lead us to the general concept of *knowledge networks*, intended as a form of overlay – distributed in a network scenario and being an integral part of the overall infrastructure – in which all the information about the context is properly represented, organised, and correlated, and around which semantically-enriched stigmergic interactions among the components of the autonomic infrastructure can take place [Par97]. That is, a distributed knowledge infrastructure representing a sort of nervous system for the autonomic communication system, across which all information and stimuli needed for the coordinated functioning of the system flow and get organized.

This position paper aims at unfolding the idea of knowledge networks and it is organised as follows. Section 2 details on the need for knowledge networks, and tries to identify what actual role and position they could assume in future autonomic communication scenarios. Section 3 elaborates on the potentials of knowledge networks in future scenarios, also with the help of a few examples. Section 4 sketches a rough research agenda and discusses related work in the area. Section 5 concludes.

2. Knowledge Networks

We are now witnessing an age of computing ubiquity where our work and home environments are increasingly enveloped by computing resources. This comes at a cost, which is the significant problem of configuration and complexity of these resources. If computing power is to serve us, and the converse is to be denied, then these resources and their rich panoply of services must be able to carry out their increasingly complex functions without significant intrusion into our lives.

These services, with underlying technology network entities encompassing autonomic computing and communication systems, require a high degree of contextual knowledge, including knowledge about the social, computational, and physical environments in which they are situated, as well as self-knowledge about their own functioning. There is a requirement for future autonomic networks to provide meaningful knowledge-based decision making, and ultimately to infuse pervasive systems and improve our human experience of interaction. This is what Weiser [Wei91] describes as the notion of *calm*, where the computing resources quietly modify themselves to suit the needs of the user.

2.1 Why Are They Needed?

Autonomic communications networks (both the network resources and the application components and services exploiting them) need to reason about their situation and to understand their own behaviour. To do this they are required (both at the level of individual components and as a whole) to be introspective and reflective, and to feed back the results of these processes to be used to improve performance. This is the *raison d'etre* to make networks smarter, to make them more self-aware, and to provide the knowledge with which they can manage themselves. In order to manage themselves, the network and its entities and services need some form of "*knowledge networks*" through which all available knowledge is properly represented, correlated, and accessed. The reasons that lead to that concept of knowledge networks are synthesised below.

Firstly, there is a basic need for expressive and flexible means to promote context-awareness. Networks, their entities and services need to have an awareness of situations with differing degrees of granularity [Ste05]. There is a requirement for some form of computational model of context processing as in [Bal00] that orchestrates context stimuli and components in a coherent representation. We also need some way to gauge the quality of our contextual information objectively as it is gathered, as from the Quality of Context mechanism of Buchholz *et al.* [Buc03], in which any contextual information comes associated with parameters including precision of information, correctness probability, trust worthiness, resolution and regency. Simply said, contextual information cannot reduce to a trivial set of data to be accessed by components, but requires some higher-form of organization.

Secondly, contextual information cannot be simply considered as local and locally available to components and services. For a satisfactory adaptive orchestration of distributed activities (whether this is intended to be the orchestrated configuration of network components or the coordination of distributed service components), the exploitation of local knowledge only may not be enough. Nor can one think of concentrating in a single site or of replicating anywhere all available knowledge, especially when this knowledge represents dynamically evolving situations, i.e., it is subject to obsolescence. The compromise solution is to enable components which need more than simply local knowledge to organise and correlate distributed knowledge into sorts of networks that enable distributed components to "navigate" through the available knowledge to attain, on demand, the required degree of contextual awareness.

Third, there is a recognised need for future autonomic communication scenarios to promote cross-layer interactions [SAC05]. This means that the service level and the network level cannot work as separated universes, each towards its own goals. Rather, a continuous exchange of information must occur between the service and the network level, and vice-versa, so as to ensure that the overall activities of the system, at each level, will contribute towards the achievement of a satisfactory functioning. For this coordination and exchange of information to occur without significant interoperability issues, there must be some place where common information can be stored and can be properly organised so as to be accessible and understandable by both the network and the application levels, and accordingly to the means proper of each level.

Fourthly, it is known that a reasonable and effective way to promote self-organization and self-adaptation (i.e., autonomic behaviour) in distributed systems is via stigmergy, i.e., by indirect interactions occurring via a computational environment in which components can spread and sense information [Par97]. The presence of a distributed network of knowledge, to be accessed for sensing and effecting by both network and application level components, can act as the computational environment to enforce stigmergic self-organization. Moreover, if such space other than simple digital pheromones can contain properly represented and correlated situational knowledge, one can think at leveraging stigmergy to more sophisticated forms of cognitive self-organization.

2.2 What form could Knowledge Networks take?

Knowledge networks are *reflective spaces* for autonomic communication systems. Being capable of storing distributed, heterogeneous, dynamically constructed, sophisticated knowledge, they can form a conceptual middle layer across which network components as well as application-level components can access information and can coordinate with each other. They act as a form of network memory, in which knowledge may be replenished continually as the network and its entities evolve and reflect introspectively. But what form do knowledge networks take, and where does the knowledge reside?

The schematic in Figure 1 illustrates how we conceive knowledge networks as a conceptual layer, positioned between the physical network level (there included the physical level, reified in the forms of the environmental information that can be produced by sensors) and the application level (there included the social level, reified in the form of social information produced by social/user profilers).

In general, the knowledge generated by both levels reaches the same conceptual knowledge level, and here it is properly put in context of extant knowledge. This means that the knowledge has to be:

- Dynamically generated and represented in proper ontological relations;
- Properly correlated, i.e., networked with existing knowledge on the basis of what it represents and of what use it and related knowledge may be to the application or the network level.

We assume that each entity in the network, whether a software agent or a network component, has the capability of accessing the knowledge network layer for reading the knowledge in it, understanding the ontological relations between different pieces of knowledge and navigating links that relate distributed knowledge. By this, components can also properly understand where newly produced knowledge can be inserted in a knowledge network, and how this has to relate with existing knowledge. That is, components have the capability of dynamically shaping the knowledge networks to have it always reflect the current overall situation of the network.

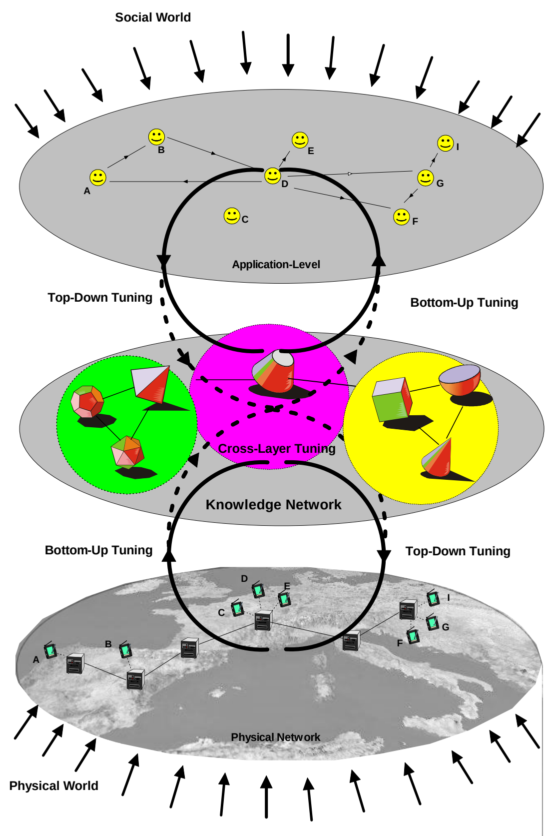

Figure 1. Knowledge Networks in Context

In general, we do not consider the presence of specific computational entities in charge of maintaining and updating knowledge networks [Cla03]. Such a solution would be too heavyweight to be general-purpose, and would introduce additional complexities. Rather, we consider that components at both the application and the network level will be directly in charge of populating, storing, and maintaining portions of fully distributed knowledge networks, arguably with the help of reactive code fragments associated to knowledge pieces and aimed at automating their update and maintenance upon changing conditions.

For the actual production and update of the network of relations in a knowledge network, we may consider an ontological construct at both network and application level that is replenished continuously via the introspective process described earlier. This behavioural feedback loop, in essence, is the knowledge generator that dynamically populates the ontology. The ontological construct must be designed to be very flexible, even to the extent of facilitating self-revision [Hef01]. It must also be capable of fusing contexts [May03] and knowledge from different ontologies as networks and devices interconnect in an *ad hoc* manner.

Given the above considerations, the knowledge networks can act as the mean via which intra-level tuning of activities may occur (see Figure 1). However, they can also act as the mean via which the application-level can tune its activities to reflect events occurring at the network-level, and vice-versa. In addition, given that the activities of a component may reflect in some change/update in the knowledge network, that some other components can sense and for which its own behaviour can be affected, the possibility of stigmergic interaction is intrinsically promoted.

In general, we consider the possibility of a multiplicity of knowledge networks to co-exist in the same overall network infrastructure, each possibly serving different application-level or network-level goals. However, the need for achieving effective cross-layer interactions and globally coherent activities may require different knowledge networks to be somehow related to each other. In particular we envision the possibility of identifying conceptually easy, practical, and scalable ways by which to compose and relate a variety of diverse knowledge networks and the diverse knowledge they contain. Specifically, we consider as promising the possibility of enforcing the construction of scale-free knowledge networks, exhibiting a self-similar structure that can facilitate robust navigation and update. Also, this could promote nesting of various knowledge networks into each other and, accordingly, could tolerate an exploitation of knowledge networks at different scales (zooming in and out depending on needs).

3. Putting Knowledge Networks to Work

The knowledge network concept outlined in this paper has the potential to make possible a sophisticated degree of autonomic behaviour in future networks by providing them with introspective, cross-layer knowledge. In this position paper, we do not have the clear visibility to help describe some proof-of-concept implementation, nor do we already have crystallised ideas about how all the above ideas could be realised. Nevertheless, we can try sketch some potential applications of the concept.

3.1 Resource Management and Load Balancing

Any distributed network infrastructures with dynamically changing resource demands requires some sort of resource management tools to have its all resources effectively exploited and to provide reasonable quality of service to the application level. For the sake of simplicity, let us focus on the load balancing issue.

Traditional distributed load balancing tools consider the presence of system level processes, devoted to handle specific local resources (whether computing, communication, or memory resources) capable of monitoring the current load on local hardware resources [ShiKS92, Xu95]. Whenever they perceive specific resources are overloaded (or under utilized), these processes engage other processes in some sort of negotiation, aimed at re-distributing the load on the system resources (e.g., by re-allocating some application-level process or by establishing different routing paths).

A variety of strategies have been proposed for distributed load balancing. Different strategies can be conceived for having local processes understand if they are overloaded or underloaded: this can rely on static non-adaptive load thresholds, or they can be based on some sorts of load information exchange with other nodes to comparatively estimate the local load. Different strategies can also be conceived for negotiation, depending on which nodes (overloaded or underloaded) initiate negotiations, and on which nodes in the system (all nodes or a limited number of "close" nodes) have to be involved in it.

However, for all the above traditional approaches, the strategy rely on local system processes to perceive local load information, possibly acquire more global information by requesting it to colleague processes on different nodes, and act on the basis of this information. Nothing is traditionally said about the possibility of organising distributed load information to promote more informed decisions without having processes to explicitly coordinate with each other every time a decision has to be taken. Nothing is traditionally said about the possibility of exploiting application-level information to enforce load re-distribution patterns that, other than satisfying the hardware viewpoint, can also accommodate specific application-level needs.

The idea of knowledge networks lets us envision a radically different approach to load balancing. Rather than having processes elaborate local information, we could think of having local load information be injected (and updated upon significant changes in value) in a knowledge network to contribute to the dynamic formation of a distributed "load field", representing in a sort of virtual landscape of the distribution of load over the network. The local value of load field and its local gradients can then be perceived by system level processes to understand "where" in the network/landscape load increases or decreases, and to somehow understand not only what is the local load, but also how such local load relates to the overall load in the network. Also, such fields can be enriched with semantic information describing e.g., the types of resources involved and any additional resource-specific information to could serve the load balancing purpose.

Given the availability of the load field, one can think of having load distribution occur by simply imposing load (i.e., the entities that actually produce such load, such as data packets for communication load and application processes for computational load) to distribute in the network by "rolling down" the load field to reach underloaded zone. This eventually achieves a satisfying (sub-optimal) balance of resource exploitation, without involving any negotiation among system-level load balancer processes. Also, provided that the load field is promptly updated upon any significant change in the load of some resources (which can be achieved via simple reactive code fragments associated to information in the knowledge network), the resulting dynamic distribution of load is made self-adaptive, in that any allocation of load to resources will automatically reflect the current global load situation. To some extent, we can consider this way of achieving load balancing as a sort of stigmergic interaction occurring via the load field.

From the application viewpoint, the above approach makes it possible for application components to play an active role in load distribution, other than the passive role of being distributed here and there. Firstly, since they too have access to the knowledge network, they can somehow “bias” the structure of the knowledge network to have it reflect their own needs. For instance, they can artificially “heighten” the shape of the load field in some zones to be ensured that they will have a specified amount of resources devoted to their execution without having other application-level components roll down to these zones. Secondly, they can enrich the load field with any type of application-specific knowledge, to be connected (via proper ontological construct) to the available load and resource information available at the lower level. In this way, one can put to work fruitful cross-layer interactions, where: on the one hand, application-level components can fruitfully exploit both types of information towards the achievement of their application goals; on the other hand, network-level components can direct application-level components towards those part of the system where their needs can be better satisfied without negatively affecting the overall systems functioning, which the availability of a semantically enriched load field enable to effectively evaluate in an introspective way.

3.2 Pervasive Computing

An application scenario which can strongly take advantage of our knowledge networks approach is pervasive computing, here intended as the support of individual and collaborative human activities in an environment densely populated by embedded computers (e.g., sensor networks and computer-based cameras), computer-enriched objects (e.g., smart furniture), and personal computing systems.

Such pervasive computing scenarios are typically open and dynamic: new computers join the scenario at any time (as carried on by humans getting in the environment or brought in via computer-enriched objects) the same as some can leave or being dismissed. The need of exploiting at the best all the available computing resources requires spontaneous inter-operability, i.e., the capability of all computing-based devices to be found in the environment (*a priori* unaware of each other and never explicitly configured to work together) to start interacting with each other towards the achievement of some application goals. Also, the scenario intrinsically involves situated computational and communication activities, in that the distributed computing infrastructure is put to service for improving humans interactions with the surrounding environment and with the current “situation” of the environment.

To solve the above problems, the pervasive computing research community recognises that middleware infrastructures based on active spaces are necessary [Rom02]. These considerations about interactions in a pervasive computing scenario occur *via* kinds of shared memory spaces, where to store and by which to access all information about the context/situation, properly organised by the space itself according to shared ontologies, and in which uncoupled (data-mediated) interactions among components (application-level and network-level) may occur without having components to know each other.

Our concept of knowledge networks leverage the active spaces approach by suggesting organising contextual/situational information into fully distributed networks of knowledge. As in active spaces, networks of knowledge can be used for uncoupled interactions among components. However, the knowledge networks approach provides for a more dynamic and lightweight perspective, in that it does not suggest that the organization of knowledge should take place by specific processes devoted to this, but rather suggest an approach in which all components contributed to the building and maintaining of the available situational knowledge. Also, it provides a better support for distributed self-organizing and self-adaptive activities, being a fully distributed knowledge network on which to rely for effective stigmergic coordination.

4. Research Agenda and Related Work

For our idea of knowledge networks idea to become a practical approach, several open research problems have be unfolded. This section analyses some of the most relevant issues – defining a broad research agenda for knowledge networks researches – and discuss how related research thrusts can somewhat contribute to it.

In general, these issues can be all generally related to the following problems (Figure 2), each of which analysed in the following sub-sections: *(i)* how to represent knowledge using proper ontological constructs; *(ii)* how to generate, compose, and relate distributed knowledge; *(iii)* how to have knowledge networks evolve and according to which structure;*(iv)* how to exploit this knowledge to achieve autonomic behaviour at both the network and the application levels.

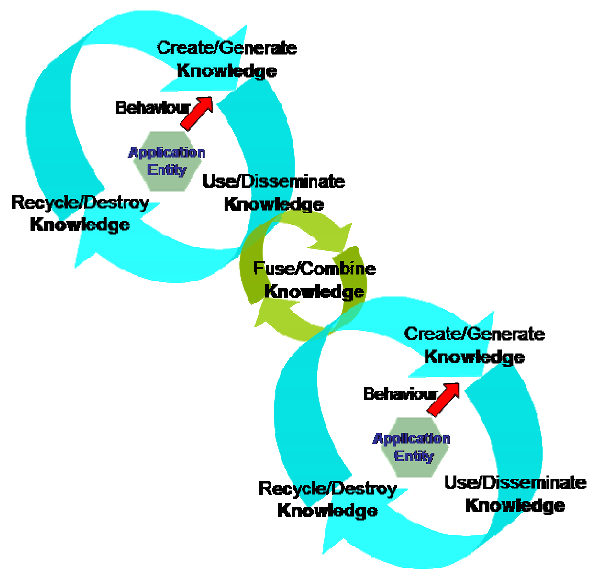

Figure 2. Knowledge Network Lifecycle

Before continuing, we emphasise that our approach here is clearly distinguished from the 'knowledge plane' approach [Cla03], and thus introduces different research issues. The knowledge plane approach considers an additional network layer between the network and the application layer, as the place in which nearly all network control activities take place. The knowledge plane is populated by heavyweight intelligent agents [ZamJW03], managing and exchanging knowledge about the current state of the network, and that directly enact forms of control over both network and application components. In our idea, instead, knowledge networks are not intended to be populated, handled, and managed, by additional knowledge-level components. Rather, to avoid the burden of an additional distributed computational layer, and to more fruitfully promote cross-layer interactions, we consider knowledge networks as managed by existing components at the application and network levels (at least supported by some simple reactive code fragments). Thus, while the research issues in the knowledge plane approach relate to how have agents in the knowledge plan interact with each other to properly control the network, the research issues in our knowledge approach relates to how components can generate, maintain, and exploit knowledge.

4.1 Defining Ontologies for Knowledge Networks

Ontological constructs [Usc96] can enable the modelling of contextual information semantically. They provide a general model which is independent of programming language, underlying operating system or middleware. Other knowledge 'consumers' in the network must be able to access and use the ontological formalisms developed. Accessing information stored in a network of distributed contextual knowledge requires the specification of information locators, e.g. in the form of an addressing scheme as well as request routing procedures. The relation between knowledge representation and addressing scheme (i.e. how can information be mapped deterministically or probabilistically to locators) as well as request routing schemas are important aspects.

One approach within the ontology category has been proposed as the Aspect-Scale-Context Information (ASC) model [Str03a]. In this model, using ontologies provides an uniform way to specify the models core concepts as well as an arbitrary amount of sub-concepts and facts, together enabling contextual knowledge sharing and reuse in an ubiquitous computing system [DeB03]. These implementations build up the core of a Context Ontology Language (CoOL), which is supplemented by integration elements such as scheme extensions for Web Services and others [Str03b].

The CONON context modelling approach by Wang *et al.* [Wan04] is based on the same idea of the ASC/CoOL approach, namely to develop a context model based on ontologies because of its knowledge sharing, logic inferencing and knowledge reuse capabilities. Wang *et al.* created an upper ontology which captures general features of basic contextual entities and a collection of domain specific ontologies and their features in each sub-domain. The CANON ontologies are serialized in OWL-DL which has a semantic equivalence to well researched description logics. This allows for consistency checking and contextual reasoning using inference engines developed for concept languages.

A promising emerging context modelling approach based on ontologies is the CoBrA system [Che03]. This system provides a set of ontological concepts to characterize entities such as persons, places or several other kinds of objects within their contexts.

All the above, may be of help to characterize the ontological constructs to be put at work for the production of knowledge in the context of autonomic communication scenarios.

4.2 Building Knowledge Network Ecosystems

How may situated knowledge networks be put to use, and how can knowledge be combined/split in differing scales of use by network entities and entity aggregates? There is need to experiment with ensembles of knowledge to facilitate knowledge consumption and use at all differing scales in our networks.

Knowledge is generated from the behaviour and behavioural analysis of individual and aggregated autonomic network- and application-level entities. This behavioural knowledge floods into an ontological construct at entity scale. We need a mechanism that distributes our behavioural knowledge dynamically. To retrieve particular knowledge, it must be possible to address information. Instead of using a fixed structure e.g., by assigning unique identifiers, we can experiment with path languages for topologies taking the 'semantic proximity' of the information in account. The goal is to be able to give directions in a fuzzy way (e.g., "follow this street, turn left on the second traffic light and walk till you see the red building"), which still offers promise in yielding an accurate and unambiguous addressing schema.

These issues have some relations with overlay networks in P2P computing [RowD01, Bab02, Rat02, Bab02, And04]. Indeed, autonomic knowledge networks will be sorts of overlays. However – unlike traditional overlays approaches in P2P computing – knowledge networks are not intended to simply support navigation of data and messages in a dynamic network of components. Rather they are intended to provide components with a local representation of the situation, that can then be used by them to adapt their behaviour e.g. to enforce properties of self-preservation, self-aggregation, and self-organisation in general. With this regard, some recent proposal for semantic overlay networks may be of great relevance [Loe04]. Semantic overlay networks are created by network nodes in P2P systems using content metrics to relate entities. Network queries are routed via the semantic overlay network, reducing the load on nodes with non-related content. Semantic overlay clusters, cluster P2P super-peers by their characteristics, enhancing search and integration significantly. Although guided by policies defined by human experts, this approach shows merit in flood reduction in overlay networks, with potential of application to overlay knowledge networks and especially knowledge network research. Thus, the study of semantic overlays may be of some relevance for the finalization of our knowledge networks concept.

Some additional source of inspiration for knowledge networks could come from some modern middleware proposals for mobile and ubiquitous computing, which consider exploiting forms of distributed data structures – to be dynamically built and self-adapting – to act as the basic mean via which adaptive coordination activities can be promoted. Such middleware proposals include among others LIME [Pic01] and TOTA [MamZ04], Smart Messages [Bor04], Limbo [Dav02]. These approaches, by having distributed data structures typically represent some application-level knowledge, definitely shares something with our knowledge networks approach. However, so far, very little has been said on the possibility of building scalable global distributed data structures in accord with some semantic relations and ontological constructs.

4.3 Enforcing Self-Similarity and Robustness

Three considerations must be made when thinking at the possible structure of knowledge networks:

- *(i)* they should somehow reflect the structure of those networks whose work they are intended to support, i.e., the application/social networks and the technological networks;
- *(ii)* they must evolve over time in an adaptive way yet preserving their properties; and

(iii) they must be scalable and promote composability.

These three issues, though, are strictly related with each other.

Both social (and application-level) networks and technological networks (e.g., the Internet and the Web) tend to evolve towards “scale-free” topologies [AlbB02]. These classes of networks, also found in biological and physical systems, exhibit neither completely random nor completely regular connection topologies [Wat98]. They are characterised by the small-world phenomenon [Mil67]; highly clustered like regular lattices, yet preserving small characteristic path lengths. Dynamical systems models with small world coupling display enhanced signal propagation speed, computational power and synchronisability, properties which can be of great importance for the effective propagation of knowledge in autonomic communication scenarios. In addition, the scale-free characteristic tends to enforce robustness and scalability in the network structure: the same overall structure is preserved as the network evolves over time; and the network exhibit the same structural properties at different levels of observation. Again, these properties would be very important for representing evolving distributed knowledge in a robust way, and for enabling a scalable way with which to structure and compose knowledge.

In summary, it will be interesting to explore how to structure knowledge networks into scale-free structures, so as to reflect the structure of the social and technological networks they support, to support robust adaptive evolution, and to support scalability and scale-free composability at different scales of observation, and to analyse the implications of this structuring.

4.4 Promoting Cognitive Stigmergy

Swarm intelligence approaches consider that global self-organizing and self-adapting behaviour can be made emerge in systems of a large number of lightweight agents that indirectly interact via the mediation of an environment [Par97, Bon99, ParB04]. Agents, by depositing and by sensing “pheromones”, and by having the environment properly diffuse pheromones according to specific laws, can – to most extent unconsciously – self-organize their global activities into robust and adaptive patterns.

Our concept of autonomic network knowledge could potentially act as a form of computational environment via which indirect, stigmergic interactions, may take place to promote self-organization and self-adaptation of activities. Still, this requires leveraging the traditional concept of stigmergy into a concept of *cognitive stigmergy*. Self-organizing and self-adaptive coordinated activities at both the network and the application level should be enforced not simply by reacting to a local concentration of meaningless pheromones. Rather, they should be driven by the actual meaning of the knowledge represented within knowledge networks.

Clearly, to preserve the advantages of swarm intelligence approaches, this should occur without requiring ants to become heavyweight agents, and a proper trade-off between the purely reactive behaviours promoted by traditional stigmergy and the purely cognitive behaviour promoted by artificial intelligence approaches have to be found.

Similar considerations can be made for those approaches to self-organisation based on indirect interactions such as the morphogen gradients of amorphous computing [Nag02, MagM04] and the field-diffusion in teams of mobile robots [McL04].

5. Conclusions

The ambitious goals of this position paper and of the associated research road map focus on the development of sophisticated knowledge representational schema for next-generation autonomic networks. Our research should deliver knowledge representation schemes and ontologies for situated and autonomic communication-intensive services, structural mapping of knowledge ensembles to network and aggregated network entities, software interfaces for programming interaction with knowledge networks, and tools, metrics and algorithms for the evaluation and monitoring of knowledge networks.

We acknowledge that the scale of research outlined in this paper is very large and that work on developing mediated network knowledge requires us to address significant *stages* of challenges. In addition to those outlined already in the paper, challenges include managing the ontology lifecycle, in particular automated knowledge acquisition for dynamic ontology construction, the use of knowledge-level techniques to address provable, correctness-preserving transformations and adaptive algorithms, and working to understand the role of planning knowledge, including understanding and changing global and local goals. It is also important to consider that the protection of use of sensitive security and privacy information raised by applying such a shared knowledge space to a highly distributed application is addressed at the design stage of a research programme such as this.

We still have no stable ideas about how these knowledge networks will look, and to which extent they will be effectively able to deliver the promise of acting as the nervous system of a future autonomic communication infrastructure. Nevertheless, this appears indeed a challenging and fascinating research topic, involving a number of related research issues likely to impact on future autonomic communication scenarios and worthy of investigation. Bringing together network and knowledge engineering to address the problems in pervasive computing shows promise, and should open up new research directions, in particular once we begin to design and implement for real-world issues using this paradigm.

Acknowledgments: Work supported by the Italian MIUR in the context of the "Progetto Strategico ICT: IS-MANET: Infrastructures for Mobile Ad-Hoc Networks", and by UK Invest Northern Ireland Networking Support using EU Structural Funds. All support is gratefully acknowledged.

References

[And04] S. Androutsellis-Theotokis, D. Spinellis, "A Survey of P2P Content Distribution Techniques", *ACM Computing Surveys*, 36(4):335-371, Dec. 2004.

[AlbB02] R. Albert, A. Barabasi, "Statistical Mechanics of Complex Networks", *Reviews of. Modern. Physics*, 74(47), 2002.

[Bab02] O. Babaoglu, H. Meling, A. Montresor, "Anthill: a Framework for the Development of Agent-Based Peer-to-Peer Systems", *Proceedings of the 22nd IEEE Conference on Distributed Computing Systems*, Vienna (A), May 2002.

[Bal00] C. Balkenius, J. Moren, "A Computational Model of Context Processing", *6th International Conference on the Simulation of Adaptive Behaviour*. The MIT Press, 2000.

[Bon99] E. Bonabeau, M. Dorigo, G. Theraulaz, *Swarm Intelligence*, Oxford University Press, 1999.

[Bor04] C. Borcea, “Spatial Programming Using Smart Messages: Design and Implementation”, *24th Int.l Conference on Distributed Computing Systems*, Tokio (J), May 2004.

[Buc03] Buchholz, T., A. Kupper, and M. Schiffers, ”Quality of context: What it is and why we need it”, *Workshop of the HP OpenView University Association 2003 (HPOVUA 2003)*, Geneve (CH), 2003..

[Che03] Chen, H., Finin, T., Joshi, A. “Using OWL in a Pervasive Computing Broker”. *Proceedings of Workshop on Ontologies in Open Agent Systems (AAMAS 2003)*, Melbourne (AU), July 2003.

[ChoP03] T. Choudhury, A. Pentland, “Modeling Face-to-Face Communication Using the Sociometer”, *ACM Conference on Ubiquitous Computing*, Seattle, WA, USA, 2003.

[Cla03] D. Clark et al., "A Knowledge Plane for the Internet", *Proceedings of the 2003 ACM SIGCOMM Conference*, Karlsruhe (D), ACM Press, 2003.

[Dav02] N. Davies, et al, “L2imbo: A distributed systems platform for mobile computing”, *ACM Mobile Networks and Applications*, 3(2):143-156.

[DeB03] J. De Bruijn, “Using Ontologies - Enabling Knowledge Sharing and Reuse on the Semantic Web”, *Tech. Rep. Technical Report DERI-2003-10-29*, Digital Enterprise Research Institute (DERI), Austria, Oct. 2003.

[DiM04] G. Di Marzo, A. Karageorgos, O. Rana, F. Zambonelli (Eds.), *Engineering Self-organizing Systems: Nature Inspired Approaches to Software Engineering*, LNCS No. 2977, Springer Verlag, May 2004.

[Est02] D. Estrin, D. Culler, K. Pister, G. Sukjatme, “Connecting the Physical World with Pervasive Networks”, *IEEE Pervasive Computing*, 1(1):59-69, 2002.

[Gel02] H.W. Gellersen, A. Schmidt, M. Beigl, “Multi-Sensor Context-Awareness in Mobile Devices and Smart Artefacts”, *Mobile Networks and Applications*, 7(5): 341-351, Oct. 2002.

[Hef01] Heflin, J., *Towards the Semantic Web: Knowledge Representation in a Dynamic, Distributed Environment*, PhD., Thesis, University of Maryland, College Park, 2001.

[KepC03] J. Kephart, D. Chess, "The Vision of Autonomic Computing", *IEEE Computer*, 36(1), 2003.

[Loe04] A. Loeser, F. Naumann, W. Siberski, W. Nejdl, U. Thaden, “Semantic Overlay Clusters within Super-Peer Networks”, *International Workshop on Databases, Information Systems and Peer-to-Peer Computing (DBISP2P)*, Berlin (D), 2003.

[MamZ04] M. Mamei, F. Zambonelli, “Programming Pervasive and Mobile Computing Applications with the TOTA Middleware”, *2nd IEEE Conference on Pervasive Computing and Communications*, Orlando (FL), IEEE CS Press, March 2004.

[May03] R. Mayrhofer, H. Radi, A.Ferscha. Recognizing and predicting context by learning from user behavior. 2003: Austrian Computer Society (OCG).

[MikM04] M. Mikic-Rakic, N. Medvidovic, "Support for Disconnected Operation via Architectural Self-Reconfiguration", *International Conference on Autonomic Computing*, New York, NY, USA, 2004.

[Mil67] Milgram, S., The Small World Problem. *Psychology Today*, 1967. 2: p. 60-67.

[McL04] J. McLurkin, J. Smith, "Distributed Algorithms for Dispersion in Indoor Environments using a Swarm of Autonomous Mobile Robots", *Proceedings of the 7th International Symposium on Distributed Autonomous Robotic Systems*, Toulouse (F), 2004.

[Nag02] R. Nagpal, "Programmable Self-Assembly Using Biologically-Inspired Multi-agent Control", *1st Int.l Conference on Autonomous Agents and Multi-agent Systems,* Bologna (I), July 2002.

[NagM04] R. Nagpal, M. Mamei, "Engineering Amorphous Computing Systems", in *Methodologies and Software Engineering for Agent Systems: the Handbook of Agent-Oriented Software Engineering*, Kluwer Academic Publishing (New York, NY), 2004.

[Par97] V. Parunak, "Go to the Ant: Engineering Principles from Natural Agent Systems", *Annals of Operations Research*, 75:69-101, 1997.

[Par04] V. Parunak, S. Brueckner, J. Sauter, "Digital *Pheromones for Coordination of Unmanned Vehicles", Workshop on Environments for Multi-agent Systems (E4MAS),* LNAI 3374, Springer Verlag, 2004.

[Pen05] A. Pentland, "Socially-Aware Computation and Communication", *IEEE Computer*, 38(3):33-40, March 2005.

[Pic01] G. P. Picco, A. L. Murphy, G. C. Roman, "LIME: a Middleware for Logical and Physical Mobility", *22nd IEEE Intl. Conference Distributed Computing Systems*, 2001.

[PicM04] G. P. Picco, A. L. Murphy, "Using Coordination Middleware for Location-Aware Computing: A Lime Case Study", *Proceedings of the 6th International Conference on Coordination Models and Languages, LNCS No. 2949,* Feb. 2004.

[Rat02] S. Ratsanamy et al., "GHT: A Geographic Hash Table for Data-Centric Storage", *1st ACM Int.l Workshop on Wireless Sensor Networks and Applications*, Atlanta, Georgia, USA, September 2002.

[Rom02] M. Roman et al., " Gaia : A Middleware Infrastructure for Active Spaces", *IEEE Pervasive Computing*, 1(4):74-83, Oct.-Dec. 2002.

[RowD01] A. Rowstron, P. Druschel, "Pastry: Scalable, Decentralized Object Location and Routing for Large-Scale Peer-to-Peer Systems", *18th IFIP/ACM Conference on Distributed Systems Platforms*, Heidelberg (D), Nov. 2001.

[Str03a] Strang, T. Service Interoperability in Ubiquitous Computing Environments, *PhD thesis, Ludwig Maximilians University Munich*, Oct. 2003.

[Str03b] Strang, T., Linnhoff-Popien, C., Frank,K. CoOL: A Context Ontology Language to enable Contextual Interoperability. In LNCS 2893: *Proceedings of 4th IFIP WG 6.1 International Conference on Distributed Applications and Interoperable Systems (DAIS2003)*, J.-B. Stefani, I. Dameure, and D. Hagimont, Eds., vol. 2893 of Lecture Notes in Computer Science (LNCS), Springer Verlag, pp. 236–247, 2003.

[SAC05] EU IST Commission, "Situated and Autonomic Communication Initiative", *Future and Emerging Technologies Report*, www.cordis.lu/ist/fet/comms.html, 2005.

[ShiKS92] N. G. Shivaratri, P. Krueger, M. Singhal, "Load Distributing for Locally Distributed System", *IEEE Computer*, 25(12):33-44, Dec. 1992.

[Ste05] Sterritt, R., M.D. Mulvenna, and A. Lawrynowicz, "Dynamic and Contextualised Behavioural Knowledge in Autonomic Communications", *Proceedings of the 1st IFIP Workshop on Autonomic Communications*. Berlin: Springer-Verlag, 2004.

[Usc96] Uschold, M. and M. Grueninger, Ontologies: Principles, methods, and applications. Knowledge Engineering Review, 1996. 11(2): p. 93–155.

[Wan04] Wang, X. H., Zhang, D.Q., Gu, T., Pung, H.K., Ontology Based Context Modeling and Reasoning using OWL. *In Workshop Proceedings of the 2nd IEEE Conference on*

Pervasive Computing and Communications (PerCom2004), (Orlando, Fl.,USA, March 2004), pp. 18–22.

[Wat98] Watts, D., J., Strogatz, S.H., Collective Dynamics of 'Small-World' Networks. *Nature*, 1998. 393: p. 440-442.

[Wei91] M. Weiser. "The Computer for the 21st Century", *Scientific American*, 265(3):94–104, September 1991.

[Xu95] M. Xu, B. Monien, R. Luling, F. C. M. Lau, "Nearest Neighbour Algorithms for Load Balancing in Parallel Computers", *Concurrency: Practice and Experience*, 7(7):707-736, Oct. 1995.

[Zam05] F. Zambonelli, M.P. Gleizes, R. Tolksdorf, M. Mamei, "Spray Computers: Explorations in Self-organization", *Journal of Pervasive and Mobile Computing*, 1(1):1-20, May 2005.

[ZamJW03] F. Zambonelli, N. Jennings, M. Wooldridge, "Developing Multiagent Systems: the Gaia Methodology", *ACM Transactions on Software Engineering and Methodology*, 12(3):410-470, July 2003.