Towards Halos Networks

Ubiquitous Networking and Computing at the Edge

Antonio Manzalini
Strategy, Future Centre,
Innovative Architectures
Telecom Italia, Italy

Télécom SudParis,
Institut Mines-Télécom
Evry, France

Noel Crespi
Télécom SudParis,
Institut Mines-Télécom
Evry, France

Vânia Gonçalves
IBBT-SMIT, Vrije
Universiteit Brussel,
Belgium

Roberto Minerva
Strategy, Future Centre,
Innovative Architectures
Telecom Italia, Italy

Télécom SudParis,
Institut Mines-Télécom
Evry, France

***Abstract*— This paper presents Halos Networks as an architectural paradigm for developing ubiquitous networking and computing services at the edge of the networks A Halos Network is like a wireless network spontaneously emerging through the interactions of distributed resources embedding wireless communication capabilities. Halos Networks are capable of delivering services and data virally through multiple devices, machines and objects interconnected with each other. This paper discusses business roles for Stakeholders in Halos Networks scenarios, specifically concluding with some remarks about the Operators' role.**

Keywords-component; Ubiquitous Networks and Services; Self-Management; Autonomic; Internet of Things.

I. Introduction

Networks are becoming more and more pervasive and dynamic, capable of interconnecting large numbers of nodes, IT resources, machines and Consumers' electronics devices embedding communication capabilities.

In the future, anything will be a network node. Actually, with the deployment of Internet of (and with) Things and Machine-to-Machine, current estimates [1] show that in a few years there will be many billions of electronic devices connected with each other and to the Internet.

In this scenario, it is realistic to imagine services and data virally delivered through multiple devices, machines and objects interconnected by a dynamic network of networks. Reasonably, this evolution will occur first at the edges of the networks, where we will see the proliferation of Sensor Networks, Personal Area Networks (PAN), Vehicular Area Networks (VAN), and in general, networks of networks and all types of machines and embedded systems.

This evolution raises technical challenges and important socio-economic issues for stakeholders to consider: from simplifying such emerging complexity when managing future networks to identifying new business opportunities and models.

This paper presents Halos Networks as an architectural paradigm for developing ubiquitous networking and computing services at the edge of the networks A Halos Network is like a wireless network spontaneously emerging through the interactions of distributed resources embedding wireless communication capabilities. For example, a halo could be the set of sensors and actuators plus a controller, a tiny PC and a smartphone creating a User's Wireless Personal Area Network (WPANs). A WPAN is a network centered around an individual person's workspace or context.

The outline of the paper is the following. In Section II an example of Halos Networks scenario is presented. Section III makes a brief summary of the state-of-the-art of autonomic networking and computing which are considered the most relevant avenues of research impacting Halos Networks development. Section IV and V focuses, respectively, on the Halos Network theoretical models, architecture and prototyping developments. Section VI discusses a number of business models for Stakeholders, concluding with a brief description of the Operators' role in Halos Networks scenarios. Finally, Section VII gives conclusions and discusses future work.

II. Halos Networks' Scenario

Consider a scenario of "ultra-dense networks" at the edge where a huge number of communicating entities, with storage and processing capabilities, are interacting locally with each other. Imagine, each person surrounded by a sort of halo (e.g. a sort of WPAN), within which devices, sensors, actuators, controllers, tiny PCs, etc. are interconnected with each other and can be autonomically "plugged and played" (with no need for human configuration).

Each device within the halo is able to communicate with its peers. The introduction of autonomic and learning capabilities is dramatically simplifying required configurations by the Users. When people's halos come into close proximity (within a few meters), the resources of the halos can flock together and spontaneously create a network of network, which is what we call Halos Networks.

Imagine that communication, storage and processing resources are clustered in "halos" not only centered around people but also Vehicles, Street Lamps, Kiosks, and so on. Actually some recent advances in ICT have literally been transforming cars into small data centers and mobile nodes of future networks [2]. Then Halos Networks will achieve wide and dynamic distributions by all halos flocking.

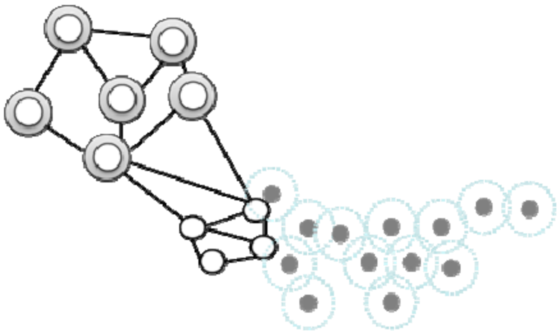

Figure 1. Halos Networks emerging at the edge

It should be noted that Halos Networks go beyond the concept of Wireless Ad-hoc networks (e.g. Mobile Ad-hoc NETworkS, Wireless Meshed Networks, Wireless Sensor Networks) [3]: actually they allow accessing (and providing) service by using and sharing local processing and storage resources. Halos Networks will look like a distributed complex communication fabric composed of large numbers of autonomic halos. This fabric is adaptive as, through the "chaos" of halos interactions and their local adaptations, collective self-organization properties emerge.

Halos Networks are fleeting but also persistent, since many objects are fixed, and others are dynamically moving. Short-to-middle range connectivity emerges through local device-to-device communications, but long-range interactions are enabled only by entering the big Net.

Moreover, different Halos Networks are learning each other's resources, functionalities and services and how to optimally make use of them. It is possible to indicate which networks can make use of the services offered. This includes support for relaying services, thereby offering services from other networks to their own neighbors.

Eventually, the progressive disappearance of Halos Networks' boundaries will pave the way towards resource symbiosis [4] over large scales.

III. State of the Art

Autonomic networking and computing are considered the most relevant avenues of research impacting Halos Networks development. There is an impressive number of publications and initiatives investigating these issues, most of which relate to architectures and component models, offering the basic building blocks with which to create autonomic self-* behaviors. This section presents a brief overview of these works.

IBM, as part of its autonomic computing initiative [5], has outlined the need for current service providers to enforce adaptability, self-configuration, self-optimization, and self-healing, via service (and server) architectures revolving around feedback loops and advanced adaptation/optimization techniques. Driven by such a vision, a variety of architectural frameworks based on "self-regulating" autonomic components have been proposed [6], [7], [8] based on the common underlying idea to couple service components with software components called "autonomic managers" in charge of regulating the functional and non-functional activities of components.

In Autonomia framework [9], the autonomic behavior of a system and its individual applications is handled by so-called mobile agents. Each mobile agent is responsible for monitoring a particular behavior of the system and for reacting to the changes accordingly.

A slightly different approach is provided by the AutoMate framework [10]. Similar to Autonomia, autonomic behavior in the AutoMate framework is handled by the agents and is implemented in the form of first order logic rules. Agents continuously process these rules and policies among themselves and perform the desired actions.

In [11], FOCALE architecture is based on mapping business level system constraints down to low-level process constraints in an approach called policy continuum [12]. This policy-based approach for specifying autonomic system behavior allows network administrators to specify business level policies for network management (using natural language), for example, defining different internet connection bandwidth rates for different users, SLA, QoS policies etc.

In [13], the Autonomic Communication Element (ACE) model is described. ACEs can autonomously enter, execute in, and leave the ACE execution environment. In general, the behavior of autonomic elements is typically provided in relation to the high-level policies that define the element's original behavior [14]. Within the ACE model, such policies (called plans) are specified through a number of states, along with the transitions that lead the ACE execution process from one state to another. Plans distinguish between the ACE's "regular" behavior, which is its behavior when no events undermining the ordinary execution occur, and the "special cases" that can occur during the plan execution process and which could affect the regular ACE execution process. If such occurrences are foreseen, the ACE behavior can be enhanced with rule modification specifying the circumstances under which the original behavior can be relinquished, along with the new behavioral directions to follow.

A very similar endeavor also characterizes several research efforts in the area of Multi-Agent Systems [15]. Multi-agents represent (de facto) the types of autonomic components which are capable of self-regulating their activities in accordance with some specific individual goal(s) and, by cooperating and coordinating with each other, according to some global application goal. However, it is worth emphasizing that that Multi-Agent Systems does not imply an autonomic behavior per-se. At the level of internal structure, Belief Desire Intention (BDI) agent systems, as implemented in agent programming systems like Jadex, JACK or Jason or in the context of the Cortex project [16], propose the use of intelligent agents to deal with autonomic and context-aware components. At the core of this model there is a rule-based engine acting on the basis of an internal component state that is explicitly represented by means of facts and rules [17], [18]. At the level of multi-agent systems and their interactions, agents are generally expected to discover each other via specific agent-discovery services, and are supposed to be able to interact.

Common to most of the proposed approaches (both those based on autonomic components and multi-agent systems) is the existence of a traditional middleware substrate to implement discovery and interactions between components or agents. On the other hand, none of the above approaches seems to address the problem of globally re-thinking ubiquitous networks as complex environment with emerging properties.

Overall, this paper is providing novel contributions by proposing a basic architectural model for policy-based, autonomic self-management of resources ubiquitously distributed in Halos Networks. The vision is that Halos Networks can flock together and create large communication network collecting many individuals that form large organized communities, where services can be spread virally. In this sense Halos Networks are leveraging the Spines [19] approach, which is a generic overlay network that provides transparent multi-hop unicast, multicast, and any-cast communication with a variety of link and end-to-end protocols.

IV. THEORETICAL MODEL AND ARCHITECTURE

As mentioned, the flocking of Halos Networks is covering the interactions of large numbers of nodes, devices, and smart objects. This complexity requires that each halos' to be both sensitive to the context variations and capable of reacting "autonomically" in order to self-adapt.

In other words, a halo should be capable of managing a set of heterogeneous autonomic resources. It can be seen as a DES (Distributed Event System): i.e. a dynamic system whose states are time-evolving as events occur. From a theoretical viewpoint, many approaches have been proposed to model DESs, most notably finite state machines, Petri nets and generalized semi-Markov processes. Among these models, Finite State Machines (FSM) represents a computation model that is the most straightforward means to control the stability behavior. FSM provides for a good understanding of the predictable problems such as controllability, observability, co-observability, normality, decentralization, and non-determinism and is the model we adopt for halos networks.

Therefore a halo can be modeled as a network of interacting FSMs, since such a network of FSMs would still be an FSM composed of k components' FSMs interacting with each other. Interestingly, non-determinism in FSMs is represented by a *choice of states* where the optimal action is yet to be decided and where it can be learned, with reinforcement learning.

Figure 2 shows the architecture of an halo consisting of a set of services. This architecture is leveraging the concept of Self-Managed Cell architecture reported in [20]. For example, the discovery service discovers resources and components being part of the halo and the other halos entering in the communication range (each single halo is clearly designed for interactions). The policy service is in charge of managing the policies specifying the halo behavior. A publish/subscribe event bus is used for interaction between halos' components and for distributing events triggering policies.

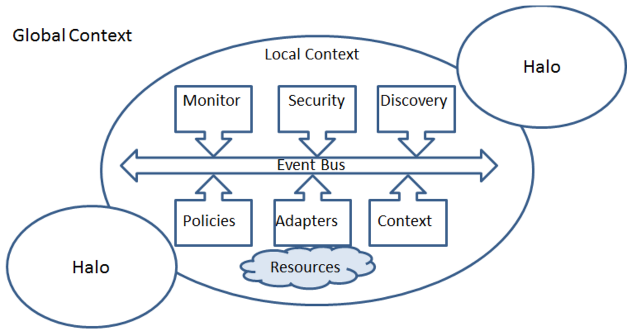

Figure 2. Architecture of a Halo

The overall architecture a Halos Network is structured into three layers, in charge of actuating three different kinds of behavior:

- Automatic behavior: this layer is in charge of fast pre-defined reactions for self-adaptation to predefined contexts and can be designed by means of automatic control-loops modeled with deterministic FSMs;
- Autonomic behavior: this layer is responsible for local adaptation achieved by exploiting halos' learning capabilities. The layer can be designed with ensembles of deterministic and non-deterministic FSMs and reinforcement learning methods;
- Globally Self-Organized behavior: this layer is in charge of diffusing local context information to orchestrate local reactions (activation-deactivation of rules) for reaching global goals (self-organization). This layer exploits a sort of "controlled" reaction-diffusion process of context information.

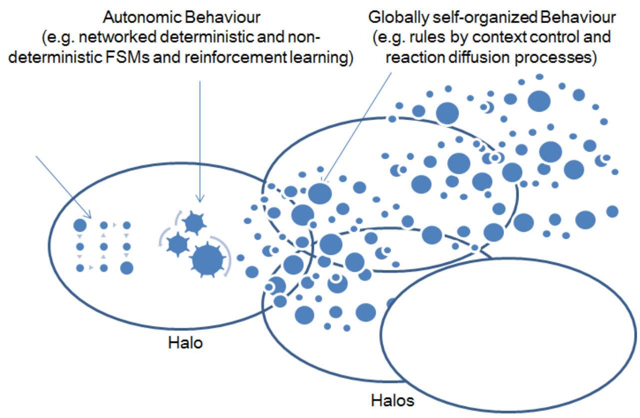

Figure 3. Overall architecture: three levels of behavior

V. PROTOTYPING

In a real proof-of-concept (under development) an halo can be easily implemented with a smart phone (acting as a Wi-Fi

Hot Spot), one (or more) cheap, tiny PCs (e.g. a Raspberry Pi costing $ 25) and one (or more) microcontrollers (e.g. based on Arduino).

The prototype is based on the architecture reported in figure 2 and 3. Specifically, each halo is empowered with the perception of its local context, i.e. the environment inside and around. Each halo diffuses its context information (e.g. tuples) hop-by-hop according to a set of propagation rules (determining how context information should be distributed and propagated in the Halos Network). Any context information can be accessed locally, simultaneously accounting for the influences of the information propagated from other halos (including the fixed ones).

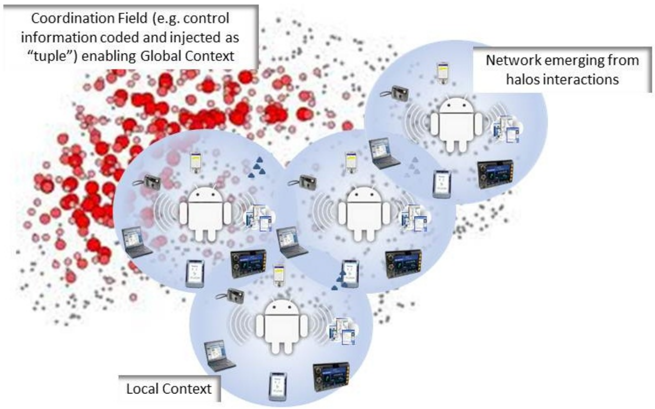

Figure 4. Halos Networks: local reactions and global self-organization

The idea is developing a sort of global coordination field (i.e. a global context) (figure 4), injected by halos (and potentially control points) in the network and autonomously propagating. In other words, halos are interacting with each other and with the environment by simply generating, receiving and propagating distributed data structures (e.g. tuples), representing context information. This field is providing halos with a global representation of the situation of the overall Halo Network (to which they belong). This coordination field is immediately usable: an halo is moving in this field like an object is moving in a "gravitational" field. Environmental dynamics and halos' local decisions will determine changes in the field, closing a feedback cycle. This process enables a distributed overall self-organization.

As mentioned coordination field can be made of tuples of data which can be injected and diffused by each node. Local reading of these tuples of data (e.g. through pattern matching) can trigger local self-adaptation behaviors. A simple event-based engine, monitoring configurations and the arrival of new tuples, reacts either by triggering propagation of other tuples or by generating events.

In fact, a number of open source applications are available on the web to implement node primitives and local autonomic behaviors.

VI. Socio Economic Issues and Business Models

During the last few years we have witnessed a diffused "consumerization" of ICT and mobile devices. This is partly a consequence of a decrease of prices of devices and Users becoming more technologically-savvy. At the same time, a whole range of applications running on top of all sort of devices has come into place enabling convenience and flexibility for Users. Seen the combination of these advantages and the gained technological knowledge, Users have become themselves producers of applications, the so-called "Prosumers". Halos Networks discussed in this paper will allow Consumers to create their "own networks" with little costs and efforts for using and sharing connectivity, storage and processing resources. Obviously, although this can become a reality in a short time, the scenario does not represent an alternative to traditional communications access services, but rather a novel and complementary approach. Consumers would still need to be connected directly, or through multi-hops, to Network Operators' infrastructure, still providing broadband access in order to reach the big Network.

In this section we intend to discuss some socio-economic issues, business models, concluding with a brief description of the Operators' role in Halos Networks scenarios.

A. Business Model Configurations

Although a wide range of business model configurations could be envisioned in Halos networks scenarios, this paper is analyzing just a few of them (those which are felt the most strategically relevant). Moreover, particular attention has been focused on the User centric Model and the role of the Network Operator.

1) User centric Model

In this model, Users have communication capable terminals. These terminals could be based on short range connectivity and optionally long term connectivity (and hence being associated to a mobile Operator). Terminals are not locked or controlled by a manufacturer (as for instance Apple), they are capable of installing open version of software. Additionally they could be capable of adopting open source firmware (much as modems or access gateways can install "openwrt", or mobile terminals can adopt an "openmoko" approach). In this configuration there is no limitation in connectivity capabilities supported by a node. The user can decide to use its nodes for supporting his halo communication needs and can also offer communication, processing and storage capabilities to nearby nodes. Such a model is extremely interesting because these open nodes could be using communication capabilities in a specific area like Wi-Fi connection, Bluetooth or other short range ones. In addition if the terminal is enabled for long reach communication, it could became an hub for other nodes. Furthermore, these kind of terminals could be able to use technologies derived from Cognitive Radio and Software Defined Radio and create very capable and powerful networks with considerable communication capabilities (e.g., whitespaces).

This model breaks the current status quo and could also be a disruptive factor for established value networks and business models. On the other hand, the practical applicability of this approach has great obstacles like current regulation, established business models and the reluctance of many Users to engage into unstable and disruptive technologies.

2) Device manufacturer centric model

In this model the Device Manufacturer takes control of not only the manufacturing and marketing roles but also controls the resource sharing provision. A potential strategy the Device Manufacturer could pursue is to restrict access to the resources that can be shared in the devices or to only give access using specific applications provided and sold by the Device Manufacturer. For the latter case, the Device Manufacturer would also take the role of application design and could sell these applications through third-party marketplaces or establishing its own marketplace.

3) Device manufacturer partnership with Network Operator model

This model is similar to the previous one since the Device Manufacturer restricts access to resource sharing but would also implement firmware or operating system functionalities to restrict Internet sharing. This would potentially allow Network Operators to have the final say in how the device could be used and get a share of resources consumption.

This model can be compared to Apple and Microsoft's implementation of tethering on smartphones, i.e., sharing a data connection through WiFi or Bluetooth with other devices. In Apple's case tethering is controlled by the Network Operator through the iOS implementation of the carrier profile specification. The carrier profile allows a mobile operator to among others restrict or give access to tethering and specifically charge for tethering traffic.

Moreover, the Network Operator could also pursue another type of partnership with the Device Manufacturer. That being of a device being sold with a data/communications subscription from a specific Network Operator. For example, all connected Toyota cars sold in the U.S. would require a Sprint data subscription to benefit from ITS functionalities.

4) Network Operator centric model

In this model, the Network Operator controls the most important roles related to device final usage and resources sharing. The Network Operator takes up the active role of marketing devices to the customer and takes control of additional functionalities. An example of this configuration is the bundle of a phone with a certain subscription or SIM-locking a phone to a specific NO. Additionally, the Network Operator could also look for added revenues by taking the role of application design and selling device or network-specific applications through third-party marketplaces or on its own marketplace.

B. Further elaborations on the role of the Operator

Currently Operators are coping with the continuous decrease in the value of bandwidth. In the future this phenomenon could slow down, but it is very unlikely that it will revert. Connectivity will be cheaper and cheaper. This will offer a lot of opportunities in terms of edge services. Terminals will become more and more intelligent and capable, and their range of connectivity possibility will increase mixing up long and short or very short-range connection capabilities. Terminals will became fully fledged communication nodes. The Operators can try to maintain the current approaches and business models, but already there is a move of the value from the networks to the terminals. Services are provided at the edge of the "network" and this trend will consolidate. Operators are at risk of losing connectivity revenues as well as service. A traditional approach, i.e., the one of providing network services, is not giving the expected results. Pursuing a conservative approach on the long run could be detrimental for Operators. In this paper another approach for the Operators is advocated: the enablers of open environments. Terminal Manufactures are trying to create walled gardens in which the terminal providers are offering to users the terminal, the services and sometimes also constraints on which operator to choose. Operators should oppose to this by enabling the openness of communication and computation environments. Open environments have a greater need to communicates (services and terminals are not confined to a specific market) and hence they give value to connectivity. This connectivity will be more and more short range, but in many cases long range connectivity is a value because it allows to connect far away halos nets and users or to access to information that is not locally available. Connectivity is a sort of disruptive factor for walled gardens as proposed by major Terminal Vendors. Operators should emphasize the value offered by open systems and connectivity, i.e., the freedom to access to any information at any time without constraints posed by the walled gardens. This value proposition is difficult to present to users that up to now are attracted by nice user experiences and very populated application stores. One consequence of walled gardens is that the behaviour of users is more observable in these confined environment and users pay a rewarding user experience in terms of a restriction in their freedom and in a strict profiling. The Operators should adverse this and enable a different environment in which people are free to cross borders, to enter and exit from environment, to be empowered and being the owners of their profiles and data. Operators should be on the side of users and promote open environments instead of following the walled garden approach of vendors.

Operators could also become infrastructure providers of Halos Networks. The sensors, the poles, the buoys needed for communicating between the user nodes and the environment could be in a large part deployed by Operators in conjunction with the Public Administrations. Deploying Networks that do support self-organization of nodes makes the life easier form a network management perspective and, as mentioned in this paper, managing a plethora of small communication nodes is not possible with traditional means and approaches. These nodes could also be organized in such a way to be able to support locally services that the Operator wants to provide to local communities or dynamic group of users. The infrastructural costs could be reduced because the used platform is the one comprised of the thousands of nodes available locally.

In summary, Operators could have a double advantage from Halos Networks: to promote connectivity by enabling open gardens and to reduce the management burden of edge nodes by promoting self-organization.

VII. Conclusions and Future Work

This paper has presented Halos Networks as an architectural paradigm for developing ubiquitous networking and computing services at the edge of the networks A Halos Network is like a wireless network spontaneously emerging through the interactions of distributed resources embedding wireless communication capabilities. Halos Networks are capable of delivering services and data virally through multiple devices, machines and objects interconnected with each other.

Preliminary proof-of-concept activities have shown that an halo can be easily implemented with a smart phone (acting as a Wi-Fi Hot Spot), one (or more) cheap, tiny PCs (e.g. a Raspberry Pi) and one (or more) microcontrollers (e.g. based on Arduino).

The paper has also elaborated on a set of business models for Stakeholders, concluding with a brief description of the Operators' role in Halos Networks scenarios. Specifically it is argued that Network Operators could have a double advantage from Halos Networks: to promote connectivity by enabling open gardens and to reduce the management burden of edge nodes by promoting self-organization.

Future work will progress the prototyping activities in order to provide also experimental results on Halos Networks. From a theoretical point of view, future work will consider: 1) potential extensions of the Shannon capacity formula for multi-source ultra-dense wireless networks; 2) investigating new Wi-Fi medium access protocols scaling better in multi-hop networks. In fact, as known theoretically, the capacity of a multi-hop wireless network increases with node density and node mobility in spite of the apparently effect of transmission interference (e.g. proportionally to the square root of the network size, i.e. number of nodes). This promises enormous wireless capacity for ultra-dense Halos Networks.

Acknoledgments

Authors wish to thank Roberto Saracco for inspiring and supporting this activity at the Future Centre of Telecom Italia.

References

[1] OECD (2012), "Machine-to-Machine Communications: Connecting Billions of Devices", OECD Digital Economy Papers, No. 192, OECD Publishing http://dx.doi.org/10.1787/5k9gsh2gp043-en

[2] http://www.siemens.com/innovation/en/news/2012/e_inno_1207_1.htm

[3] A. Goldsmith et alii, "Beyond Shannon: The Quest for Fundamental Performance Limits of Wireless Ad Hoc Networks" - IEEE Communications Magazine, May 2011;

[4] R. Gedge, Symbiotic networks, BT Technology Journal, 21(3), 2003, pp. 67–73;

[5] Kephart J. & Chess D.M.: The Vision of Autonomic Computing. In: IEEE Computer, vol. 36 no. 1 pp. 41-50, IEEE Computer Society Press;

[6] White, S. R., Hanson, J. E., Whalley, I., Chess, D. M., & Kephart, J. O.: An Architectural Approach to Autonomic ComputingInternational Conference on Autonomic Computing;

[7] Liu, H., & Parashar, M.: Component-based Programming Model for Autonomic Applications. International Conference on Autonomic Computing, pp. 10-17;

[8] Xu, J., Zhao, M., Fortes, J., Carpenter, R., & Yousif, M.: On the Use of Fuzzy Modeling in Virtualized Data Center Management. 4th International Conference on Autonomic Computing, Jacksonville (FL), USA, 2007;

[9] Dong, X., Hariri, S., Xue, L., Chen, H., Zhang, M., Pavuluri, S., & Rao, S.: Autonomia: an autonomic computing environment. IEEE international conference on Performance, Computing, and Communications;

[10] Agarwal, M., Bhat, V., Matossian, V., Putty, V., Schmidt, C., Zhang, G., Zhen, L., Parashar, M., Khargharia, M,. & Hariri, S.: AutoMate: enabling autonomic applications on the grid. Autonomic Computing Workshop, pp. 48-57;

[11] Jennings, B., Van der Meer, S., Balasubramaniam, S., Botvich, D., Foghlu, M.O., Donnelly, W., & Strassner, J.: Towards Autonomic Management of Communications Network. IEEE Communications Magazine, vol. 45 no. 10 pp. 112-121;

[12] Van der Meer, S., Davy, A., Davy, S., Carroll, R., Jennings, B., & Strassner, J.: Autonomic Networking: Prototype Implementation of the Policy Continuum. 1st Workshop on Broadband Convergence Networks;

[13] Manzalini, A., Zambonelli, F., Baresi, L., Di Ferdinando, A.: The CASCADAS Framework for Autonomic Communications. In A. Vasilakos, M. Parashar, S. Karnouskos, & W. Pedrycz (Eds.): Autonomic Communication, Springer Verlag ISBN: 978-0-387-09752-7, pp. 147-168;

[14] Parashar, M., Hariri, S.: Autonomic Computing: An Overview. In: Unconventional Programming Paradigms, Springer-Verlag, ISBN: 3-540-27884-2 pp. 257-269;

[15] Valckenaers, P., Sauter, J., Sierra, C., & Rodriguez-Aguilar, J. A.: Applications and Environments for Multi-Agent Systems. Journal of Autonomous Agents and Multi-Agent Systems 14(1);

[16] Biegel, G., Cahill, V.: A Framework for Developing Mobile, Context-aware Applications. In: International Conference on Pervasive Computing and Communications, Orlando (FL), USA;

[17] Klein, C., Schmid, R., Leuxner, C., Sitou, W., & Spanfelner, B.: A Survey of Context Adaptation in Autonomic Computing. International Conference on Autonomic and Autonomous Systems, Gosier, Guadeloupe, 2008;

[18] Li, J., Powley, W., Martin, P., Wilson, K., & Craddock, C.: A Sensor-based Approach to Symptom Recognition for Autonomic Systems. International Conference on Autonomic and Autonomous Systems, Valencia, Spain, 2009;

[19] Spines overlay network: http://www.spines.org, 2010;

[20] Sloman M., Lupu E.: Engineering Policy-Based Ubiquitous. Systems Published by Oxford University Press on behalf of The British Computer Society, 2005; doi:10.1093/comjnl/bxh000.