

Guided Self-Organization for Edge ICT Fabrics

A. Manzalini¹

Abstract— This paper is arguing that technology advances and cost reductions in processing, storage and communications will determine a growing amassing of resources at the edge of current networks (i.e., in the distribution and access areas, up to end Users premises). This trend, intertwined with new emerging paradigms such as Software Defined Networks, will impact deeply the evolution future networks and services, enabling the deployment of novel flexible architectures capable of creating a galaxy of new ICT business opportunities. It is argued that the sheer number of nodes, devices and smart systems being deployed at the edge will create a sort of distributed “fabric” offering an enormous processing and storage power. On the other hand, the level of complexity and dynamism of said “fabric” will pose challenging requirements for the orchestration and management of edge resources, which could be faced by deploying a sort of bottom-up self-organization meeting classical top-down management approaches.

Keywords-component; Autonomic, Self-Organization, Edge Networks, SDN, Nfv, Standard Hardware, Future Networks.

I. INTRODUCTION

The term “fabric” has been used in the past to refer to a distributed computing system consisting of loosely coupled storage, networking and processing functions interconnected by high bandwidth links. The term has also been used to describe a flat, simple intra data centres network optimized for horizontal traffic flows, mainly based on a concept of “server-to-server connectivity”. In this paper, it is argued that the edge (i.e., in the distribution and access areas of current networks, up to end Users premises) is going to become a sort of “fabric”, i.e., a distributed platform consisting of loosely coupled processing storage and networking pervasive resources interconnected by wired and wireless link.

In particular, it is argued that technology advances (e.g., standard hardware performance, embedded communications, device miniaturization, etc.), and the related costs reductions, will determine such a large amassing of resources at the edge of current networks that, if properly managed, will be able to offer an enormous processing and storage power, at low costs.

Actually, “intelligence” has already started moving towards the edge since a few years, and it appears very probable that in the next few years the trends technology performance vs costs will allow a radical change of paradigm: from having a network of services to applications made available by a plethora of (small) Telco-ICT Players.

This will impact dramatically future networks evolution (not only at the edge, but even in the core segment), allowing not only cost-savings and QoS improvements, but even creating new business

opportunities. Actually, technology and business developments will be more and more strictly intertwined in the future. Certain technologies and solutions will be adopted not only if advantageous (e.g., reducing costs) and trusted but also if they will enable desired business ecosystems (with the related foreseen business models); on the other hand, newly designed ecosystems will look for enabling solutions and technologies capable to bring them into reality. Metaphorically, it will be like in Nature, where evolution selects the winning species: the winning ICT services will succeed, grow, and promote further investments, while losing ideas will fade away.

In this perspective, it will be strategic for current and future Players to explore the several challenges and opportunities offered by the exploitation of the edge “fabric” potentialities (e.g., in terms of QoS improvements, cost savings and new business models).

In fact, the edge domain is arguably the most active and critical segment of today's networks, in terms of innovation, strategy, and investment. Services and traffic dynamics will be more and more in the hands of Users (and by end Users it is meant not only people but also machines, smart objects, appliances and any device which is attached to the network at the edge). Nevertheless, the growing level of complexity and dynamism of said edge fabrics of resources will pose challenging requirements, from the management and control points of view, which will be covered only by deploying automatic and autonomic capabilities to enable a sort of guided self-organization in the edge fabric.

This should be seen in comparison with the current “network ossification” which is hampering innovation. In fact, traditional centralized processes (which worked perfectly in the past) are now creating a lot of limitations for the development and deployment of new network functionality, services, security designs, management policies and approaches, and other elements that are essential to cope with the increasingly challenges of future ICT networks and services.

One concrete example is represented by the number of middle-boxes [1] deployed in current networks: not only these nodes are contributing to the “network ossification”, but also they represent a significant portion of the network capital and operational expenses (e.g., due to the management effort that they require). Basically, a middle-box is a stateful node supporting a narrow specialized network functions (e.g., layer 4 to 7); it is based on purpose-built hardware (typically closed and expensive). Examples of said functions are Wide Area Network (WAN) optimizers, Network Address Translation (NAT), performance-enhancing-proxies, intrusion detection and prevention systems, any sort of firewalls, other application-specific gateways, etc. Transforming fully in software these middle-boxes would determine several advantages, such as cost savings and increased network flexibility.

Emerging paradigms such as Software Defined Network (SDN) [2] and Network function Virtualization (NfV) [3], are likely to offer the opportunity to develop, fully in software, middle-boxes' network functions, and to allocate them dynamically according to Users' needs and Providers' policies. Some of these network functions are already available today in open source software and, in principle, could be executed as applications on Virtual Machines (VM) running

¹ Manuscript received May 2013, revised February 2014. A. Manzalini is with Telecom Italia - Strategy/Future Centre, Innovative Architectures, Telecom Italia, Turin, Italy. (E-mail: antonio.manzalini@telecomitalia.it.)

standard hardware servers (as the ones available in Data Centres for provisioning Cloud Computing services). There are open source implementations of firewalls, load balancers, proxies and caches, monitoring and measurement, intrusion detection, and ubiquitous NAT. Even if today there are still some concerns about the performance of network functions fully developed in software and running on standard hardware, performance gap (with respect to purposed-built hardware solutions) is decreasing rapidly. Even today performance can be improved by utilizing hardware accelerators or by running multiple instances of the software (utilizing multiple processing elements).

As such, as an example, it is argued that will be possible allocating and executing said network functions in the edge fabrics, where an incredible amount of processing and storage resources are progressively accumulating. This will bring several techno-economic advantages, even to traditional Network and Service Providers, if they will be able to capture and adapt to the coming shift of paradigm.

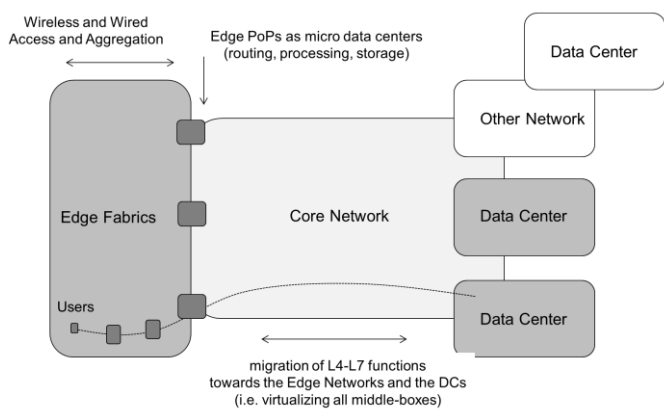


Fig. 1. Edge Fabrics

The outline of the paper is the following. In Section II provides some considerations about how approaching constrained optimization problems using the principles of statistical physics (e.g., in analogy with the symmetry breakings phenomena). Section III makes a brief summary of the prior art about autonomic systems. Section IV describes the scalable node concepts and the overall functional architecture of the Edge ICT Fabrics. Finally, Section V gives the conclusions and discusses recommended future activities.

II. CONTROL OPTIMIZATION AND SYMMETRY BREAKINGS

In the near future the edge of communications networks will be composed by such a sheer number of resources that it will look like large scale distributed systems, i.e., a sort of fabric. As known, there are strict analogies between large scale systems and statistical mechanics. In this section we'll explore how statistical mechanics and autonomic principles can help us in deploying a guided self-organization in the Edge ICT Fabric.

As well known, statistical mechanics provides a framework for relating the microscopic properties of individual atoms or molecules to the macroscopic properties of materials. In other words provides a molecular-level interpretation of macroscopic thermodynamic quantities such as work, heat, free energy, and entropy. For example the behavior of a gas (e.g., modelled by the state equation) can be described at the microscopic level in terms of the position and velocity of each molecule, which appear as random processes. Local

random behavior of the molecules causes the gas as a whole to solve a large scale constrained optimization problem.

Similarly the behaviour of electrons in an electrical network can be described in terms of random walks. This simple description at the microscopic level leads to rather sophisticated behavior at the macroscopic level: for example, patterns of potentials in a network of resistors is minimizing heat dissipation for a given level of current. Again this is like to say that the local random behavior of the electrons causes the network as a whole to solve a large scale constrained optimization problem.

Moving to small living entities, a termites colony can be seen self-organized ecosystem. We can imagine that each termite has its own utility functions (for example concerning aspects related its survivability): evolution selected those behaviors capable of maximizing said utility functions. Termites are interacting each other, and with the environment, thus cross-influencing their behaviors and having an impact on the environment. Again this is like saying that the termites' behaviors and cross-interaction allow solving a global utility i.e. large scale constrained optimization problems.

In summary, one may say any large scale (eco)systems tends to move from the state with higher energy (i.e., higher cost) to the state with lower energy (i.e., lower cost); local minima of this functional are usually related to the stable stationary states (functional minimization). Indeed, maximize profits, minimize costs, minimize the loss are typical economics problems, which can be mathematically modelled as constrained optimization problems.

It is argued that the same approach could be taken for designing and operating an highly flexible network architecture at the edge capable of self-adapting dynamically to changing conditions: the guided self-organization will become a property emerging from the interactions of nodes local behaviours, solving constrained optimization problems.

A typical approach in similar directions is looking for minimization (or maximization) of an objective function subject to constraints on the possible values of the independent variables. For example "Layering as Optimization Decomposition" [4], [5] integrates the various protocol layers into a single coherent theory, considering them as carrying out an asynchronous distributed computation over the network to implicitly solve a global Network Utility Maximization (NUM) problem.

Since then, many research studies have been carried out on distributed network resource allocation using the language of NUM. These efforts have found many applications in network resource allocation algorithms and Internet congestion control protocols, e.g., [6], [7], [8], [9]. For example, the TCP/IP protocol can be seen as an example of an optimizer: its objective is to maximize the sum of source utilities (as functions of rates) with constraints on resources. In fact, each variant of a congestion control protocol can be seen as a distributed algorithm maximizing a particular utility function. The exact shape of the utility function can be reverse-engineered from the given protocol. Similarly, other recent results also show how to reverse engineer Border Gateway Protocols (BGPs) as a solution to the Stable Path Problem, and contention-based Medium Access Control (MAC) protocols as a game-theoretic selfish utility maximization.

Also cross-layer interactions can be characterized by viewing the process of network layering as the decomposition of a given NUM problem into many sub-problems. These sub-problems are then "combined together" by certain functions of the primal and dual variables [5].

Normally maximization/minimization of complex functions (or functionals) are achieved with metaheuristics. On the other hand,

these approaches are not that scalable (and quickly converging) for highly distributed systems composed by huge amount of interacting nodes, which is the case for the edge fabrics.

This paper adopts another perspective, typical of statistical physics, which is achieving guided self-organization as a property emerging from the interactions of nodes local behaviours.

Let's imagine an edge fabric where a number of nodes are interacting and their interactions are depending on a set of control parameters. The performance of a swarm of nodes can be described by a reward associated to an average over the nodes states and over time. This is typical problem considered in statistical physics when dealing with the symmetry breakings² in large ensembles of interacting particles.

Mathematically speaking, symmetry is characterized by the invariance of some mathematical object under some transformation. For example, a parabola $y=x^2$ is symmetrical with respect to the y -axis, since it is invariant under the transformation that takes the variable x and transforms it into $-x$. In classical physics, symmetry breakings imply conservation laws: for instance, translation invariance implies momentum conservation, while rotational invariance implies angular momentum conservation. There are many systems that exhibit symmetry breakings phenomena: Bose-Einstein condensates, superfluids, superconductors, ferromagnets, anti-ferromagnets, crystals, and, according to the Standard Model, particles.

In the context of this paper the term symmetry breaking indicates the emergence, in swarms of interacting nodes, of a guided self-organization leading to solve large scale constrained optimization problems. A key question is if and how distributed control strategies can break the symmetry of the interactions in a guided way.

As a simple example, imagine introducing into the nodes automatic and autonomic features in terms sets of "controllable" local rules, making the nodes able to learn to compute input-output data mapping with desired actions or properties (local rules means rules capable of changing for example the "strength" of the interconnections of a node in the immediate neighbourhood; moreover this would be in line with the Hebb's postulate of learning).

A. Example of Symmetry Breaking in an Edge Fabric

Let's consider a wireless network composed by a large number of entities. It can be demonstrated mathematically that for a certain number of nodes, and for a critical communication range (i.e., the control parameter) the probability that the network is fully connected shows a sudden transition 0-1 (e.g., that could be seen as a phase transition). This is an example of symmetry breaking under guided with the control parameter, i.e., communication range.

The theoretical capacity of the network is proportional to the square root of the network size (number of nodes): e.g., one million nodes with available bandwidth of 1 Mb/s can reach a total capacity on order of Gb/s if the control parameter is properly tuned.

Let's imagine now that Users may wish sharing not only the wireless bandwidth but also their local storage and processing resources: the fabrics, in principle, could abruptly provide an enormous capacity. In fact, again, for a certain number of nodes, and a critical level of resource sharing, the overall throughput of the network shows a sudden transition like a symmetry breaking (in this

² This phenomena occurs if a macroscopic state of the system is not symmetric with respect to the symmetry of the interactions in the system.

case the control parameter could be the level of incentives to Users – or even to other Providers - to encourage to share their physical resources).

Several start-ups are already offering today storage and disaster resilience services using decentralized and distributed virtual resources, shared by Users. So, indeed, the concept of harnessing idle resources is starting to be extended also to the processing idle power distributed at the edge, up to the Customers' premises. Examples of provisioned services will be CDN-like services, content sharing, aggregation, transformation, data collection, etc.

On a large scale, nevertheless, this would require guided self-organization capability for properly orchestrating said local idle storage and processing resources when executing and provisioning network functions and other services.

B. Example of Use Case

A Network-Service Provider may want to provide Users' with the features that personal data (e.g., stored in edge micro Data Centres) and applications (e.g., executed by local edge resources) are following them seamlessly when they are moving from one network attachment point to another one (even when this implies the move to other fabrics crossing the Core network).

In other words, it should be possible to move data and VMs executing services seamlessly with no impact on QoS/QoE perceived by the Users. It should be possible, also, for security or other policies to follow logically specific network applications (e.g., running on VMs). These features could be extended to data and services associated to Users in order to build distributed virtual data centers at the edge (this is an ideal service, for example, for Universities, Enterprises, etc.), provided at costs which are a small fraction of traditional cloud computing services.

Let's imagine now an event (e.g., Olympic Games) where a crowd of people is converging towards a certain venue. The aggregation of a crowd of Users in space and time can be seen as a symmetry breaking event. The Network-Service Provider should make sure that the network and service fabrics is following opportunistically the same symmetry breaking by dynamically moving data and VMs to provide Users with the requested ICT services with the expected QoS. This can be achieved through a guided self-organization of the Edge ICT Fabrics.

III. PRIOR ART ON AUTONOMICS

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 and incomplete overview of these works.

IBM, as part of its autonomic computing initiative [10], 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 [11], [12], [13] 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 [14], the autonomic behaviour 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 [15]. 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 [16], FOCAL architecture is based on mapping business level system constraints down to low-level process constraints in an approach called policy continuum [17]. 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 [18], 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 behaviour [19]. 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 [20]. 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 [21], 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 [22], [23]. 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.

It should be mentioned that RTD activities on autonomic networking and self-managing networks have reached a maturity level so to start moving results toward standardization of architectural principles of the Self-Managing Future Internet. This is done in the Industry Specification Group (ISG) on Autonomic network engineering for the self-managing Future Internet (AFI), under the auspices of the European Telecommunications Standards Institute (ETSI). Specifically the AFI architectural reference model is a set of fundamental design and operational principles, describing the functions, processes, and interfaces, and so it provides Decision-making Elements (DEs), responsible for autonomic management and control of network resources [24].

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.

In [25], [31] a number of results and proposals have been presented about self-aware networks, i.e., networks capable of exploiting self-adaptiveness. Several questions have been addressed such as scaling, security, reliability and mobility.

The novelty introduced by this paper is providing a simple architectural model, based on three layers, for a guided self-organization of resources ubiquitously distributed in the edge of current networks. The vision is that edge resources can flock together and create a sort of fabric collecting many individuals that form large organized communities, where services can be spread virally.

IV. AN ARCHITECTURE FOR GUIDED SELF-ORGANIZATION

We are arguing that the edge fabrics will be composed by a sheer number of nodes devices, and smart objects which are flocking, being interconnected by wireless and wired links. The complexity, deriving from the number and dynamism of the interconnections, requires that each node of the fabric to be both sensitive to the context variations and capable of reacting “autonomically” in order to self-adapt to changing conditions.

Figure 2 is showing the concept of an edge node which could scale from a Customer Premises Equipment up to an Edge PoP. In particular, an Edge PoP can be seen as the future evolution of current edge routers, i.e., a sort of micro-data center at the border between access-aggregation (edge) and core networks.

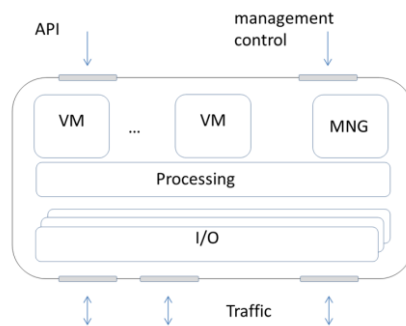


Fig. 2. Example of scalable node concept

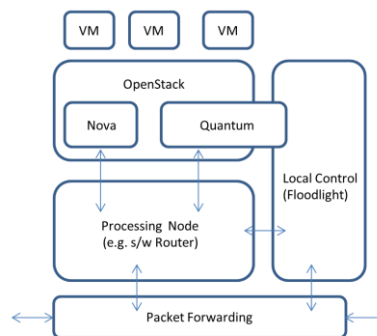


Fig. 3. Another example of node (Edge PoP) concept (with Open Source Software and Standard Hardware)

As a further example, figure 3 is showing, for example, a possible implementation of the Edge PoP, in terms of available open source software (e.g., OpenStack, Floodlight) and standards hardware (e.g., x86) for packet forwarding. This type of node could offer network

functions fully developed in software and executed in VMs. Moreover the node could be programmable offering sets of APIs related to control parameters steering towards self-organization.

In principle, these edge nodes can be seen as Distributed Event Systems (DES), 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, probably, the computation model that is the most straightforward means to control performance indicators. FSM provides for a good understanding of the predictable problems such as controllability, co-observability, normality, decentralization, and non-determinism.

Therefore each node can be modelled as a network of interacting FSMs with programmable control points. 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. This is for further study.

From a purely functional perspective, a node could be seen as consisting of a set of services, leveraging the concept of Self-Managed Cell architecture reported in [32]. For example, the discovery service discovers resources and components being part of the node and the other nodes entering in the communication range (each single node is clearly designed for interactions). The policy service is in charge of managing the policies specifying the node behaviour. A publish/subscribe event bus is used for interaction between node' components and for distributing events triggering policies.

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

- Automatic behaviour: this capability is achieved with fast pre-defined reactions for self-adaptation to predefined contexts and can be designed by means of automatic control-loops modelled with deterministic FSMs and deployed into edge nodes;
- Autonomic behaviour: this capability is responsible for local adaptation and it is achieved by exploiting unsupervised learning capabilities. The layer can be designed with ensembles of deterministic and non-deterministic FSMs and reinforcement learning methods deployed into edge nodes;
- Self-organization emergent behaviour: this is an emergent capability achieved through the orchestration of local reactions (e.g., through activation-deactivation of rules) by means of local context information diffused for reaching the global goals. It is a sort of "guided" reaction-diffusion process of context information.

In other words, self-organization emergent behaviour could be achieved by developing a sort of global coordination field (i.e. a global context), injected by nodes (and potentially control points) in the network and autonomously propagating. All nodes 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 the nodes with a global representation of the situation of the fabric (to which they belong). This coordination field is immediately usable: a node is moving in this field like an object is moving in a "gravitational" field. Environmental dynamics and nodes' 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 behaviours.

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. A number of open source applications are available on the web to implement node primitives and local autonomic behaviours.

V. CONCLUSIONS

This paper has proposed the vision whereby recent advances in standard hardware technologies and open source software are creating the conditions for exploiting highly innovative network architectures, mainly at the edge of current networks.

As a matter of fact, the sheer number of nodes, devices and systems being deployed at the edge, up to Users' premises, are offering an enormous processing and storage power. It is argued that exploiting these resources, closer to the Users, to execute network functions and services will bring several advantages, both in term of improved performance and cost savings (e.g., determined by the removal of middle-boxes). Moreover, the exploitation of these principles at the edge of current networks will transform this area into a fertile ground for the flourishing of new ICT ecosystems and business models.

The paper has also presented the concept of an edge node (based on standard hardware) which could scale from a Customer Premises Equipment up to an Edge PoP. In particular, an Edge PoP can be seen as the future evolution of current edge routers, at the border between the core and the edge fabrics. The overall functional architectures of an edge fabric has also been presented showing the exploitation of programmable automatic and autonomic capabilities, leading to the emergence of the guided self-organization.

Preliminary simulations and experiments are showing the feasibility of this vision. Next steps will continue this and will include also techno-economic analysis, aiming at simulating the impact of diverse cooperation-competition strategies in guided self-organization.

Lessons learnt, up today, is that tremendous technology advances are indeed making possible to develop L3 to L7 network functions (almost) fully in software (e.g., in Virtual Machines) and to allocate and move VMs dynamically on distributed resources. This trend is accelerating day by day. In the short-medium term (e.g., in five years) these trends are likely to create "tipping points", beyond which new equilibria and new ecosystems will emerge in the Telco-ICT.

References

- [1] IETF RFC3234 <http://tools.ietf.org/html/rfc3234>;
- [2] White paper on "Software-Defined Networking: The New Norm for Networks" <https://www.opennetworking.org/>;
- [3] White paper on "Network Functions Virtualisation"; http://portal.etsi.org/NFV/NFV_White_Paper.pdf NFV;
- [4] Chiang M., Low S. H., Calderbank A. R. & Doyle J. C. (2007). Layering As Optimization Decomposition: A Mathematical Theory of Network Architectures. *Proceedings of the IEEE*, Vol. 95, No. 1., pp. 255-312;
- [5] Palomar D.P. & Chiang M. (2006). A Tutorial on Decomposition Methods for Network Utility Maximisation. *IEEE Journal Selected Areas Communications*, vol. 24, no. 8, pp. 1439-1451;
- [6] Low S. H. (2003). A duality model of TCP and queue management algorithms. *IEEE/ACM Transactions Networking*, vol. 11, no. 4, pp. 525-536;

- [7] Kunniyur S. & Srikant R. (2003). End-to-end congestion control: utility functions, random losses and ECN marks. *IEEE/ACM Transactions Networking*, vol. 10, no. 5, pp. 689–702;
- [8] La R. J. & Anantharam V. (2002) Utility-based rate control in the internet for elastic traffic. *IEEE/ACM Transactions Networking*, vol. 9, no. 2, pp. 272–286, 2002;
- [9] Kelly F.P., Maulloo A. & Tan D. (1998). Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*; 49:237–252;
- [10] Kephart J. & Chess D.M.: The Vision of Autonomic Computing. In: *IEEE Computer*, vol. 36 no. 1 pp. 41-50, IEEE Computer Society Press;
- [11] White, S. R., Hanson, J. E., Whalley, I., Chess, D. M., & Kephart, J. O.: An Architectural Approach to Autonomic Computing International Conference on Autonomic Computing;
- [12] Liu, H., & Parashar, M.: Component-based Programming Model for Autonomic Applications. International Conference on Autonomic Computing, pp. 10-17;
- [13] 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;
- [14] 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;
- [15] 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;
- [16] 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;
- [17] 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;
- [18] 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;
- [19] Parashar, M., Hariri, S.: *Autonomic Computing: An Overview*. In: *Unconventional Programming Paradigms*, Springer-Verlag, ISBN: 3-540-27884-2 pp. 257-269;
- [20] 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);
- [21] Biegel, G., Cahill, V.: *A Framework for Developing Mobile, Context-aware Applications*. In: *International Conference on Pervasive Computing and Communications*, Orlando (FL), USA;
- [22] 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;
- [23] 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;
- [24] ETSI ISG: <http://portal.etsi.org/afi>;
- [25] E. Gelenbe, R. Lent, A. Nunez “Self-Aware networks and QoS”, *Proceedings of the IEEE*, Vol. 92 (9), pp. 1478--1489, 2004;
- [26] E. Gelenbe “Users and services in intelligent networks”, *Proceedings IEE (ITS)*, 153 (3): 213-220, 2006;
- [27] E. Gelenbe, P. Liu and J. Lainé “Genetic algorithms for route discovery”, *IEEE Trans. Systems, Man and Cybernetics B*, 36 (6): 1247-1254, 2006;
- [28] E. Gelenbe and M. Gellman “Can Routing Oscillations be Good? The Benefits of Route-switching in Self-aware Networks”, *15th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2007)*, October 24-26, 2007, Istanbul, Turkey, pp. 343-352, IEEE Computer Society;
- [29] E. Gelenbe and M. Gellman “Oscillations in a Bio-Inspired Routing Algorithm”, *IEEE 4th International Conference on Mobile Adhoc and Sensor Systems, MASS 2007*, 8-11 October 2007, Pisa, Italy, pp. 1-7, IEEE. <http://dx.doi.org/10.1109/MOBHOC.2007.4428681>;
- [30] E. Gelenbe and G. Loukas “A self-aware approach to denial of service defence”, *Computer Networks* 51 (5). pp. 1299-1314, 2007;
- [31] E. Gelenbe “Steps toward self-aware networks”, *Communications ACM*, 52(7):66-75, July 2009;
- [32] 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.



Antonio Manzalini was born in Turin, Italy. He received the M.Sc. Degree in Electronics Engineering from the Politecnico of Turin, in 1988. He joined Telecom Italia in 1990, starting RT&D activities on Optical Transport Networks. In this avenue of activities he was actively involved in leading positions in several EURESCOM and EC funded Project. He also chaired (during two consecutive Study Periods) two ITU-T Questions, in charge of ITU-T Standards on transport networks, such

as SDH and OTN. Since 2005 he has been involved in the context of international research programmes concerning the evolution towards Cloud Computing and Networking. Also in this context he was actively involved in leading positions in EC funded initiatives. In 2008 he has been awarded with the International Certification of Project Manager by PMI. He is member of IEEE, where he coordinates a Working Group on SDN in the Future Direction Committee. He acted as TPC member of several IEEE conferences and, recently, he was the General Chair of the IEEE Workshop “SDN for Future Networks and Services” (Trento, Nov. 2013). He is author of a book on Network Synchronization, several other book chapters and about 90 papers. He has been awarded with 6 patents on telco networks and ICT systems. Currently, he is Senior Manager at the Strategy Department of Telecom Italia addressing mainly RT&D activities concerning architectures, solutions and business models for Telco-ICT future networks and services (e.g., specifically dealing with Software Defined Networks, Network Function Virtualization, Autonomic/Cognitive Networking and Computing).