

Emerging Architectures for Post-IP Networking

Daniel F. Macedo¹, Aldri Luiz dos Santos², Guy Pujolle¹

¹Laboratoire d'Informatique Paris VI
Université Pierre et Marie Curie-Paris6, France

²Informatics Department
Federal University of Paraná, Brazil
{Daniel.Macedo, Guy.Pujolle}@rp.lip6.fr, aldri@inf.ufpr.br

Abstract

The increasing share of mobile multimedia devices accessing the Internet, coupled with the convergence of telephony networks into all-IP solutions, are bringing a new class of applications and scenarios to the TCP/IP architecture. Those scenarios are mostly characterized by their mobility, transmission of video and audio streams that may require strict QoS guarantees and use multiple access technologies (WiFi, WiMax, Bluetooth, etc). The use of the Internet Protocol (IP) for those devices is attractive due to the ubiquity of IP networks and due to their low cost. However, the IP architecture is built on several design decisions that complicate the support of the emerging devices that are being connected to the Internet. Hence, several architectures and/or protocol replacements to the current set of standards, which are generically called in the literature as “post-IP”, have been proposed. This paper describes the challenges and existing solutions in post-IP protocols and architectures. First, we will explore why the TCP/IP architecture is not tailored to mobile networks, briefly describing its limitations and

possible solutions. Next, we show a brief survey of the existing post-IP architectures, identifying their contributions and limitations.

1. Introduction

Due to the evolution of wireless networking equipment, nowadays it is routine to find PDAs, cellphones or laptops with one or more types of networking technologies, such as WiFi, Bluetooth, GSM and infrared. These so-called “convergent” devices allow networking anytime, anywhere [18]. All those technologies have different characteristics, such as bandwidth, range, latency and price. Further, users expect ubiquitous access to information with minimum service disruption as possible, even on the move. This problem involves the interconnection of equipment with different communication protocols and the roaming among heterogeneous networks. Roaming is necessary since it provides the means to an “always-on” service. If a user leaves the range of his office WLAN, for example, he must now route all traffic through its GSM interface in order to maintain his connections active. However, since each network uses a different set of protocols and paradigms, roaming is not trivial. Further, portable devices are used to access multimedia content, such as voice, music, video or interactive games. Those traffic types are much harder to provide than typical data streams, because they might require one or more characteristics such as high bandwidth, low latency and jitter and constant throughput.

The challenges are similar on public safety communications. In such situations, one or more teams of public service workers (e.g. firemen, police, army, medical services) combine to act on some catastrophic event. As each group is independently managed, they might employ different equipment and networking standards. Similar to commercial convergent networks, rescue operations will mostly demand multimedia and real-time traffic, such as voice communication or video feeds of severely damaged regions [42]. The biggest difference between “regular” convergent networks and communication for rescue situations is the lack of pre-deployed networking infrastructure, since it may be either destroyed, disabled or too congested. Thus, a number of small ad hoc networks, interconnected by means of a WAN backbone, will form a

mesh network [10] of mobile ad hoc networks. Due to the mission-critical nature of the work, the network must be extremely reliable, requiring minimal or no human intervention.

Finally, with the decreasing cost and miniaturization of electronic components, wireless sensor networks are becoming a reality [11]. Hence, devices such as IP television and stereo sets, household appliances and wireless sensors for home monitoring and environmental studies are more and more being connected to the Internet, in order to provide real-time information for their owners wherever they are [41]. Some of those devices have severely constrained resources, and thus cannot run IP [34], requiring either a simplified version of IP or a new addressing scheme. Also, the location of the node is more important than its unique identifier, creating the need for service and data location protocols that use geographic information to identify the most suitable data source [38]. Further, the need for extremely low power consumption frequently leads to cross-layer approaches or to radically new networking architectures [12, 25, 29, 40].

To achieve the convergent networking vision, several capabilities must be added to the existing networking infrastructure. First, in order to support user mobility, the network must implement handover protocols and policies [13]. Due to frequent changes in the access network and on the environment, the operation of the network must be context-aware, adapting automatically and transparently to changes.

Since IP networks are the *de facto* standard for communication, convergent architectures should be IP-based to facilitate development and deployment. However, the IP architecture is built on several design decisions that complicate the support of convergent networking. This occurs not because of poor design, but because the overwhelming success of IP spurred the adoption of networking on devices and scenarios that were impossible to imagine at the time of its inception. The requirements for mobile and ubiquitous access are not fulfilled by the current Internet technologies because it was developed for fixed, resourceful nodes that were always connected to the network and exported quasi-static services. The future Internet, on the other hand, will be based on mobile nodes, which will be at times disconnected from the network and will provide services dynamically, depending on their context.

Since the shortcomings of IP are rooted on some of its most important design decisions, building a new network architecture from scratch based on revised requirements would greatly simplify the conception of protocols for the networks of the future. Hence, several proposals for new networking architectures arose on the literature. Those architectures are generically termed post-IP, as they are considered to be the next evolutionary phase of the Internet. The amount of research initiatives focused on the subject such as US's FIND [3] and GENI [5], and recent European Union project calls such as FIRE [4] and FET [2] illustrate the importance of the topic on the networking community.

This paper provides an overview of convergent networking and surveys the existing proposals. Section 2 shows the limitations of IP-based networks when applied on a converged scenario. Next, Section 3 describes the most prominent architectures and projects for post-IP networking. Section 6 comments the trends on the research of post-IP architectures, and section 5 draws the conclusions and future work.

2. Limitations of IP-Based Networks

Networks are converging to all-IP solutions, and as such future networks will likely be based on IP or will use the IP infrastructure in order to profit from the hardware and software already deployed. However, the Internet Protocol was not designed to support the kinds of applications and node dynamics that users demand of convergent networks. We discuss below some shortcomings of IP-based networks when employed on mobile scenarios and show how each could be solved.

1. **IP addresses are used for routing and connection identification.** As the Internet was primarily designed for fixed networks, IP addresses are used for routing and identification at the transport layer. When we consider mobile hosts, however, addresses change as users roam through networks with different administrative domains. Since TCP flows are identified by a quadruple formed of IP addresses and ports, if a user changes his IP or the content migrates to another host, the connection is lost. Mobile-IP solves the problem of node mobility by adding servers which forward pack-

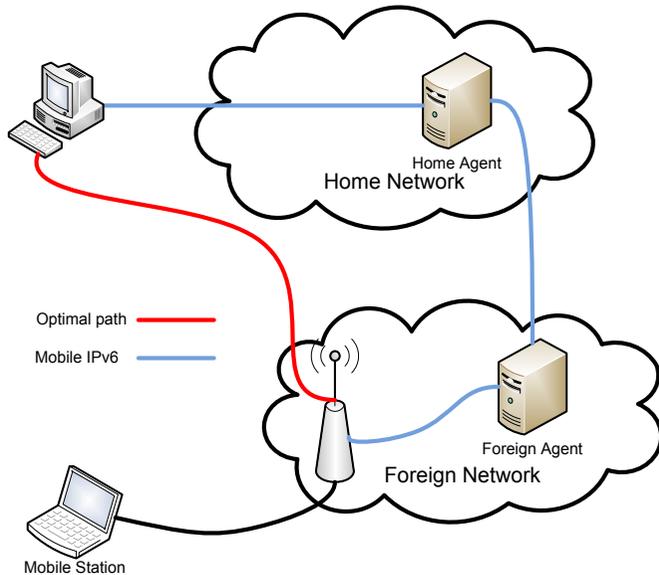


Figure 1: The traffic triangulation problem, caused by the coupling of routing and identification on TCP/IP addresses.

ets to the current location of the roaming nodes [35], as Figure 1 shows. This solution creates the triangulation problem, since data is routed to the home network, which then forwards the packets to the instantaneous address of the mobile station. This problem would not exist if naming and addressing were decoupled. Connections would be based on names, which could be dynamically reassigned to new addresses when necessary. Balakrishnan et al. go even further, proposing three levels of names in order to separate (i) users/services; (ii) hosts; and (iii) routes [15]. Thus, not only one connection can migrate from one network to another with a mobile terminal, but it can also migrate from one host to another, allowing process mobility.

2. **The intelligence of the network is usually on the edges.** The Internet is based on the End-to-End (E2E) principle, which

states that the core network must be as simple as possible, leaving the complexity to the edge nodes [37]. Although the E2E principle allowed the Internet to foster applications that could never have been imagined when the network was designed, this rule must be violated in order to allow seamless communication. Take, for instance, a user watching a video on a PDA through a WLAN. If the user goes out of range of the WLAN, the PDA might switch to a WAN network such as cellular. Some cellular networks, however, are not IP-based, thus an intermediary node is necessary to translate IP packets to the most appropriate WAN abstraction. This problem occurs even when all-IP networks are employed. For example, if a user watching a multicast video switches from WiFi access to WiMax, his available bandwidth might change. The network, in turn, could automatically transcode the video to adjust to the new bandwidth limitations. Further, even though IP is E2E, intermediary boxes that add application-aware functions to the network are recurrent on the Internet nowadays. However, the violation of the E2E principle might create incompatibilities with other functions of the network. Caches do not work in conjunction with IPSEC and SSL, while NAT hinders the use of P2P applications [19, 22].

3. **Data is usually bound to a host.** On most Internet protocols, data is associated with the machine where it is stored. FTP, HTTP and TCP are some examples. To get the file *foo.bar* on HTTP, for example, the user employs a URL such as *http://www.server.com/news/foo.bar*. This URL ties the file *foo.bar* to a server and to a directory on the machine. Thus, the file cannot be easily moved from one directory or from one machine to another. Further, for several types of information, users do not care about the source of the information, as long as they receive it. As an example, a user requests the weather information for his town from whatever meteorological station. Peer-to-Peer networks and search engines were created to bridge this gap, providing more user-friendly means to refer to the sought information. Such services arose from the difficulty of finding information on the Internet. Since those services are built over an address-based system, content-based addressing is implemented as an overlay, and

requires the implementation of an index that maps characteristics of the information sought into networking addresses.

Several solutions, such as JINI, Salutation and others were created to counter this problem [16, 17, 39]. However, since the underlying network makes requests based on addresses, service and data mobility is a big issue. Whenever a service is inserted, deleted or moved, a lengthy notification process must be undertaken. All nodes on the network that are using the service or that have cached its location must be informed of its new position. Another problem with content being bound to an address is that CDN-like mirroring requires a non-trivial set of mechanisms, such as DNS redirects and explicit content rewriting [32]. An approach that counters those limitations is content-based routing [30]. Since routes are based on meta-data, roaming services are handled seamlessly. Moreover, if multiple sources for the same data are found, the best source is selected by the routing algorithm.

As an example, if someone wishes to retrieve the latest information on the world cup, the user must somehow find out the node where the information is located and where it is logically stored. Hence, if the address or the directory where the information changes, the user will be unable to access the information. However, if queries were based on attributes, as in a search engine where the content is found using keywords, the information would be allowed to be moved around on the network transparently to the user.

4. **Management is done outside the network.** On the “classic” Internet, there are a small number of nodes to be managed, namely routers, which tend to be reasonably immutable. Thus, existing management solutions are based on operators that remotely configure the behaviour of the network. With the connection of ubiquitous devices to the network (sensors, RFID, electronic appliances) and due to mobility, services and nodes are added and removed at a much higher rate, making human-based management impossible due to scalability and complexity factors. Hence, management must be moved from the operator to the nodes, creating self-managed networks. Those networks, which are also called *au-*

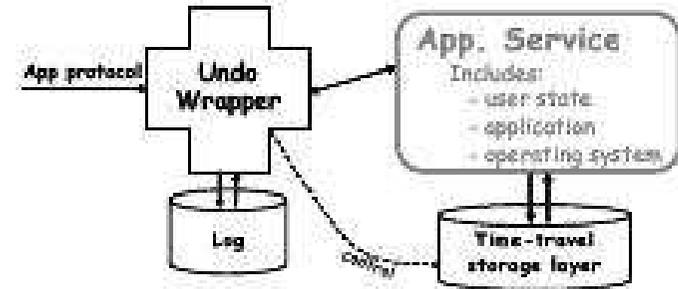


Figure 2: The generic architecture of a undo-based systems for autonomic Internet serves.

tonomic networks due to the use of the autonomic computing concept [28], employ intelligent software that monitors the conditions of the node and its surroundings, triggering configuration changes in order to avoid service degradation. In order to do that, the network must incorporate generic healing mechanism. As an example of such mechanisms, ROC (Recovery-Oriented Computing) applies the idea of *do*, *undo* and *replay* of transactional databases in generic servers to counter unknown errors [21]. As Figure 2 shows, by caching user requests and storing a known stable configuration of the server, it is possible to automatically recover from errors. Once an error is detected, the system is restored to a safe state, and the operations sent by the users are re-run.

3. The Proposed Post-IP Architectures

This section surveys the existing architectures for convergent networks. Even though there are other projects and propositions under development [1, 3, 5, 7, 33], we do not show them on this paper because their architectures are still subject to change.

3.1. The NewArch Project

The NewArch project, developed by MIT, USC and Berkeley, proposes a new architecture to the Internet, focused on the years to come [24]. Based on the assumption that the network could be redesigned from scratch, the authors devise a set of rules and principles that should guide the new architecture. These principles are implemented by the protocols FARA, RBA and XCP.

FARA implements a new architectural model that decouples routing and node identification [23]. In FARA, the functions of the IP address are broken into *entity* and *association*. An entity might be a node, thread, process, cluster and so on, while an association maps a packet to a certain entity. This flexibility allows the mobility of both ends of the flow among nodes or networks, simplifying the use of mobile agents. RBA (Role-Based Architecture) is a new approach to routing [20]. Since the Internet not always obey the principles of layering and network transparency, the authors abolish layering and introduce *roles*. Each packet is associated with one or more roles, which dictate how packets are routed and treated inside the network. Each role, thus, can be seen as a small program, which defines how packets are treated. With this abstraction, it is possible to insert “middle boxes” (firewalls, NAT, proxies) without violating the architectural principles of RBA. As on the Internet, routing is based on regions, however users might choose in which regions the packets might/might not pass, building customized paths when needed.

XCP (eXplicit Control Protocol) is a transport protocol that tackles the problems of TCP on links with high packet error rate and/or high *throughput* \times *delay* product [31]. The first situation arises because TCP cannot differentiate a packet loss due to congestion from a packet loss due to transmission problems. TCP always assumes that a packet is lost due to congestion, which is quite reasonable on wired networks. On wireless networks, TCP slows its data rate when packets are dropped, while it should keep on sending packets in the same data rate. On high *throughput* \times *delay* links, the algorithm used to determine the ideal data rate on a link increases its data rate too conservatively (additive increases). Hence, several round-trip-times are needed for a connection to reach its maximum data rate, wasting a significant amount of band-

width. XCP decouples flow control from packet retransmission, hence packets are retransmitted if an acknowledgement has not been received after a pre-determined interval, while the data rate is determined by the minimum available bandwidth on each link of the route.

The strongest point of the architecture is XCP, a transport protocol that was demonstrated to be much more efficient than TCP for a wide range of scenarios. The rest of the architecture was not evaluated, due to the difficulty of making experiments nowadays on new architectures. The architecture is quite disruptive, due to the use of roles, which would require the upgrade of every router on the Internet.

3.2. Plutarch

Crowcroft et al. take the opposite approach of NewArch. The disparate capacities of networked devices, ranging from limited sensor networks to super-computers, preclude the development of one single solution. Thus, as it is not possible to devise a new architecture that fits all conceivable scenarios, the authors build an adaptation architecture, called Plutarch, allowing the communication between different networking paradigms [26].

The Plutarch architecture interconnects heterogeneous networks, allowing each domain to use the protocols that best fit their needs. Gateways link networks with different *contexts* (protocols or architectures). Gateways implement *interstitial functions*, mapping addresses and names from one context to another. The architecture loosely defines the basic operations that must be implemented on the gateways. Those were designed to be simple and generic, to allow Plutarch to operate over any network. In Plutarch, gateways become a key part of the architecture, as they are responsible for translating data from one networking paradigm to the other. Thus, all the capabilities of the architecture lie on the efficient implementation of the gateways.

Plutarch is a very simple proposal that deals only with the interconnection of networks. Since each network may operate differently, Plutarch does not address QoS or management or security issues. However, the network allows each sub-network to implement the most appropriate solution to QoS. This does not guarantee end-to-end QoS, which must be negotiated between sub-networks, and is not currently

supported by Plutarch. Another problem with the architecture is that it does not address session and terminal mobility.

3.3. The AdaptNet Project

The AdaptNet project, developed at the Georgia Institute of Technology, implements what the authors call “Next-fourth generation” wireless systems, where nodes with heterogeneous access protocols exchange multimedia content [9]. One of the goals is to provide an evolution to IP-based networks, requiring only modifications on the mobile nodes, where a set of cross-layer protocols are employed. The heart of the architecture is the R²CP (Radical Reception Control Protocol) transport protocol, which automatically adapts the coding of multimedia content in order to optimize the perceived quality of the stream. A medium adaptation layer, located between MAC and routing, performs adaptive congestion control and coding.

Node mobility is addressed on the transport level. Using signal level measurements, nodes know approximately when they will roam into other networks, so they open new TCP connections through the new base station before the handover is completed. R²CP switches its transport between the open TCP connections, in order to allow a transparent handover to the application. For handling IP addresses and routing mobile packets, AdaptNet uses technologies such as Mobile-IP.

AdaptNet provides solutions to terminal mobility, however the proposal does not seem to account for the interconnection of non-TCP/IP networks with the Internet, as it occurs today with mobile phones that switch between Wi-Fi and GSM/CDMA. It does not solve the problems of identification and content-based search either, since the proposal does not involve addressing. QoS or autonomic management aspects were not taken into account.

3.4. Ambient Networks

Ambient Networks (ANs) are All-IP networks consisting of several heterogeneous networks [8, 36]. The project was launched from the vision of *networks beyond 3G*, where mobile users access multimedia content over a number of radio technologies. ANs employ a new IP-based

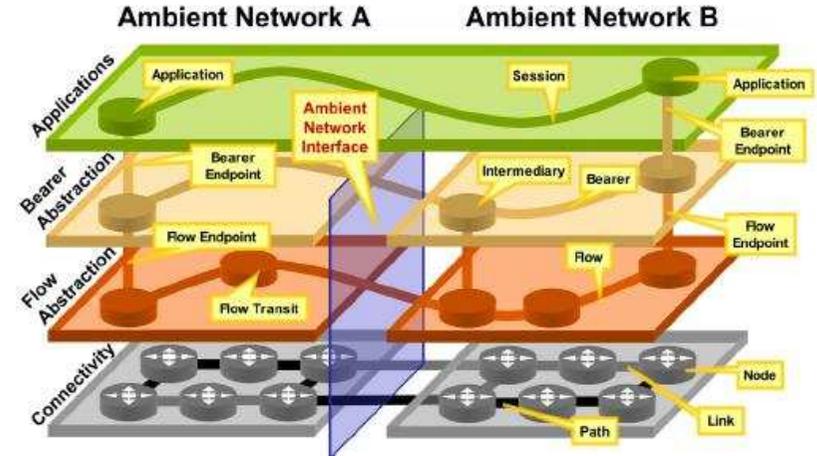


Figure 3: The abstractions used in Ambient Networks to ensure transparent mobility and service migration.

layer called *ambient control space*, which sits on top of routing. This layer solves the problems that arise from the interconnection of heterogeneous networks. All operations on the ambient control space are autonomic. Because ANs allow the composition of larger networks without human intervention, the architecture could also be used in emergency situations.

In order to provide a transparent communication to the application layer among different networking paradigms, the ambient control space employs several layers of abstraction, as depicted in Figure 3. On the border of every Ambient Network domain, special nodes, called *translators*, act as gateways, negotiating end-to-end QoS parameters and “repacking” the payload to adhere to the conventions of each network. Translators can also implement content conversion gateways or transparent filtering/caching. Translators allow each sub-network to employ the best protocols for their operation, optimizing the communication on the sub-net level, while the ambient control space takes care of end-to-end QoS provisioning. Ambient networks use Entity-based endpoints,

called *Bearer*s. Bearers are one abstraction over transport protocols, thus an application flow might traverse several connections. This flexibility allows the use of middle-boxes and seamless node and process mobility, as connections can be opened or closed on demand. Finally, since the communication is broken into pieces, each of them employs the protocols more suited to the domain.

The Ambient Network project focuses on the autonomic management and composition of networks. Thus, its main contribution is its autonomic management plane, the ambient control space. Their solution addresses mobility of terminals and applications by decoupling identification and routing. No QoS aspects were evaluated at the moment.

4. Technology Trends in Post-IP

This section comments on the design decisions of the architectures as a whole, in an effort to point the direction on the research on post-IP networking. We close this section by showing the concept of virtualisation applied to networking.

Most of the surveyed proposals lean towards the autonomic networking paradigm to manage their architectures. Still, those solutions do not specify how the autonomic management solutions will interface with the networking protocols. Security and privacy issues are not a popular topic among the surveyed architectures, even though they will be one of the cornerstones of future systems.

Naming is treated on most architectures, since it is the root cause of the limited mobility support of IP. The architectures propose the use of abstract names that are independent from the device hosting the information. Another recurrent pattern is the interconnection of several networks into one larger network using gateway nodes. On the Internet, this role was performed by IP, forcing all the layers above it to abide to a common communication paradigm. The proposed gateways are more flexible, as they do not require both networks to standardize the higher layers. Hence, simpler networks are able to communicate with complex ones without breaking the tenets of the architecture.

Another key concept in future networks is virtualization. Nowadays, routers are a very specialized piece of equipment, hence they are very

resistant to change, complicating the deployment of new services and/or network paradigms. The idea is to virtualize routers in order to replicate the operation of large-scale testbeds for network experimentation, such as PlanetLab [6] and the still in construction GENI [5], which allow the concurrent execution of several networking experiments over the same hardware. Each experiment, which uses its own networking stack, receives an amount of resources (bandwidth, CPU capabilities, etc) for a pre-determined amount of time. Virtual networks are isolated from each other, thus researchers may experiment freely without bothering other tests being run at the same time.

Anderson et al. were the first to propose an Internet architecture that follows the ideas of large scale testbeds. In such architecture, routers provide “slices” of their capabilities (e.g. bandwidth, processor, etc) to several isolated virtual networks [14], allowing the execution of different protocol stacks over the same hardware. Hence, such an architecture would allow several networking standards, implementations and versions to coexist. As an example, if an operator wishes to test a new upgrade or implement a complex service, he allocates a small amount of router resources to testing. This operation is safe, as the virtual router guarantees that errors on the experimental network will not affect the production network.

CABO (Concurrent Architectures are Better than One) goes even further, proposing the separation of hardware providers from service providers [27]. Telecommunication operators and ISPs would own routers and links, which in turn would be rented to service providers. It is up to the service provider to manage the rented network, using the protocols and configurations that best fit the needs of the service. The total cost of operation will be lower, as the cost of the hardware will be amortized by several clients that use the same hardware for their service. As an example, a WiFi provider nowadays must install access points (APs) on all its target area, even though there might be other companies that have underutilized APs on the same area. With a virtual Internet, the same provider could make agreements with companies that own spare AP capacity, to rent their unused resources.

5. Conclusion

The networking technologies of tomorrow will allow users to access their data and services wherever they are, no matter what type of device they use. Moreover, the network will provide a fair level of quality of service, even when users roam through different domains or access networks. This network will have one or two orders of magnitude more nodes than the Internet of today, due to the introduction of “things” and sensors. Since the Internet has been designed with a completely different usage in mind, existing networks are becoming ripe with patches and workarounds to counter its design limitations. Thus, a significant amount of effort has been spent on the proposition of new networking architectures that depart from the TCP/IP paradigm.

This paper provided an overview of the architectural aspects of convergent networking and a survey of the existing proposals. We argued that the TCP/IP architecture has several shortcomings when applied in convergent scenarios, as it arose out of a context where mobility and dynamicity were not an issue. A survey of the literature showed that there is not yet a proposal that stands out as the most suitable replacement for the Internet. Due to the complexity of the theme, the existent solutions tend to specialize on some of the aspects of the future convergent networks. While some architectures focalize on enhancing the QoS and providing mobility to wireless stations, others provide the interconnection of heterogeneous networks and rethink layering in order to provide a more evolvable network. Still, one common point in most architectures is the adoption of autonomic management and configuration techniques. In terms of architecture, the application of virtualisation on routers and stations will allow the deployment of service-aware networks, where the protocols will be fine-tuned to the provided service.

The field of post-IP networks is ripe with research challenges and problems. Most proposals do not take into account dependability, QoS and security aspects, even though network operators will not switch to an entirely new architecture unless it is proven to be reliable and secure. Finally, in order to validate new ideas that must work in planetary scale networks, it is necessary to build big testing facilities, using configurable nodes, where the whole protocol stack can be changed.

References

- [1] Clean Slate Design for the Internet. <http://cleanslate.stanford.edu/>.
- [2] FET - Future and Emerging Technologies. http://cordis.europa.eu/fp7/ict/programme/fet_en.html.
- [3] FIND - Future Internet Network Design. <http://find.isi.edu>.
- [4] FIRE - Future Internet Research and Experimentation. <http://cordis.europa.eu/fp7/ict/fire>.
- [5] GENI - Global Environment for Network Innovations. <http://www.geni.net>.
- [6] Planetlab – an open platform for developing, deploying and accessing planetary-scale services.
- [7] The 100x100 Clean Slate Project. <http://www.100x100network.org/>.
- [8] Bengt Ahlgren, Lars Eggert, Börje Ohlman, and Andreas Schieder. Ambient networks: bridging heterogeneous network domains. In *IEEE 16th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, volume 2, pages 937–941, September 2005.
- [9] Ian Akyildiz, Yucel Altunbasak, Faramarz Fekri, and Raghupathy Sivakumar. AdaptNet: An Adaptive Protocol Suite for the Next-Generation Wireless Internet. *IEEE Communications Magazine*, 42(3):128–136, March 2004.
- [10] Ian Akyildiz and Xudong Wang. A survey on wireless mesh networks. *IEEE Communications Magazine*, 43(9):523–530, Sept 2005.
- [11] Ian F. Akyildiz, Weilian Su, Yogesh Sankarasubramaniam, and Erdal Cayirci. A Survey on Sensor Networks. *IEEE Communications*, 40(8):102–114, 2002.

- [12] Ian F. Akyildiz, Mehmet C. Vuran, and Özgür Akan. A cross-layer protocol for wireless sensor networks. In *Annual Conference on Information Sciences and Systems*, pages 1102–1107, March 2006.
- [13] Ian F. Akyildiz, Jiang Xie, and Shantidev Mohanty. A survey of mobility management in next-generation all-ip-based wireless systems. *IEEE Wireless Communications*, 11(4):16–28, August 2004.
- [14] Thomas Anderson, Larry Peterson, and Jonathan Turner. Overcoming the internet impasse through virtualization. *IEEE Computer*, 38(4):34–41, April 2005.
- [15] Hari Balakrishnan, Karthik Lakshminarayanan, Sylvia Ratnasamy, Scott Shenker, Ion Stoica, and Michael Walfish. A layered naming architecture for the internet. In *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM)*, pages 343–352, 2004.
- [16] Matthias Baldauf, Schahram Dustdar, and Florian Rosenberg. A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing*, 2(4):263–277, 2007.
- [17] Paolo Bellavista, Antonio Corradi, Rebecca Montanari, and Alessandra Toninelli. Context-aware semantic discovery for next generation mobile systems. *IEEE Communications Magazine*, 44:62–71, September 2006.
- [18] Rober Berezdivin, Robert Breinig, and Randy Topp. Next-generation wireless communications concepts and technologies. *IEEE Communications Magazine*, 40(3):108–116, March 2002.
- [19] Robert Braden, David Clark, Scott Shenker, and John Wroclawski. Developing a next-generation internet architecture. <http://www.isi.edu/newarch/DOCUMENTS/WhitePaper.pdf>, July 2000.
- [20] Robert Braden, Ted Faber, and Mark Handley. From protocol stack to protocol heap: role-based architecture. *SIGCOMM Computer Communications Review*, 33(1):17–22, 2003.
- [21] Aaron B. Brown and David A. Patterson. Rewind, repair, replay: three r’s to dependability. In *Proceedings of the 10th workshop on ACM SIGOPS European workshop: beyond the PC*, pages 70–77, 2002.
- [22] B. Carpenter and S. Brim. Middleboxes: Taxonomy and issues, 2002.
- [23] David Clark, Robert Braden, Aaron Falk, and Venkata Pingali. FARA: reorganizing the addressing architecture. In *Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture (FDNA)*, pages 313–321, 2003.
- [24] David Clark, Karen Sollins, John Wroclawski, Dina Katabi, Joanna Kulik, Xiaowei Yang, Robert Braden, Ted Faber, Aaron Falk, Venkata Pingali, Mark Handley, and Noel Chiappa. New arch: Future generation internet architecture. Technical report, MIT/USC/University College London, December 2003. <http://www.isi.edu/newarch/iDOCS/final.finalreport.pdf>.
- [25] Marco Conti, Gaia Maselli, Giovanni Turi, and Silvia Giordano. Cross-layering in mobile ad hoc network design. *Computer*, 37(2):48–51, 2004.
- [26] Jon Crowcroft, Steven Hand, Richard Mortier, Timothy Roscoe, and Andrew Warfield. Plutarch: an argument for network pluralism. *SIGCOMM Computer Communications Review*, 33(4):258–266, 2003.
- [27] Nick Feamster, Lixin Gao, and Jennifer Rexford. How to lease the internet in your spare time. *SIGCOMM Computer Communications Review*, 37(1):61–64, 2007.
- [28] A. G. Ganek and T. A. Corbi. The dawning of the autonomic computing era. *IBM Systems Journal*, 42(1):5–18, 2003.
- [29] The NSF Wireless Mobile Planning Group. New architectures and disruptive technologies for the future internet: The wireless, mobile and sensor network perspective. Technical report, NSF Wireless

- Mobile Planning Group Workshop, August 2005. Cites cross-layer (objective 4.12) and autonomic management with augmented information as an important requirement for future networks. Tem refs para probing!
- [30] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In *Proceedings of the sixth annual international conference on Mobile computing and networking*, pages 56–67, 2000.
- [31] Dina Katabi, Mark Handley, and Charlie Rohrs. Congestion control for high bandwidth-delay product networks. In *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM)*, pages 89–102, 2002.
- [32] Balachander Krishnamurthy, Craig Wills, and Yin Zhang. On the use and performance of content distribution networks. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement (IMW)*, pages 169–182, 2001.
- [33] Antonio Manzalini and Franco Zambonelli. Towards Autonomic and Situation-Aware Communication Services: the CASCADAS Vision. In *IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications (DIS)*, pages 383–388, June 2006.
- [34] G. Montenegro, N. Kushalnagar, J. Hui, and D. Culler. Transmission of IPv6 packets over IEEE 802.15.4 networks. IETF Draft.
- [35] Andrew Myles, David B. Johnson, and Charles Perkins. A mobile host protocol supporting route optimization and authentication. *IEEE Journal on Selected Areas in Communications*, 13(5):839–849, June 1995.
- [36] Norbert Niebert, Henrik Abramowicz, Göram Malmgren, Joachim Sachs, Uwe Horn, Christian Prehofer, and Holger Karl. Ambient networks: an architecture for communication networks beyond 3G. *IEEE Wireless Communications*, 11(2):14–22, April 2004.
- [37] Jerome H. Saltzer, David P. Reed, and David D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–288, November 1984.
- [38] Karim Seada, Marco Zuniga, Ahmed Helmy, and Bhaskar Krishnamachari. Energy-efficient forwarding strategies for geographic routing in lossy wireless sensor networks. In *Proceedings of the 2nd international conference on Embedded networked sensor systems (SenSys)*, pages 108–121, 2004.
- [39] Siva Sivavakeesar, Oscar F. Gonzalez, and George Pavlou. Service discovery strategies in ubiquitous communication environments. *IEEE Communications Magazine*, 44:106–113, September 2006.
- [40] Vineet Srivastava and Mehul Motani. Cross-layer design: A survey and the road ahead. *IEEE Communications Magazine*, 43(12):112–119, December 2005.
- [41] Sandy Teger and David J. Waks. End-user perspectives on home networking. *IEEE Communications Magazine*, 40(4):114–119, 2002.
- [42] M. Voorhaen and C. Blondia (editors). WIDENS project deliverable 2.1: User requirements and first system architecture design. Technical report, Information Society Technologies, 2006. <http://www.comlab.hut.fi/projects/WIDENS/D2.1.pdf>.

Author’s Biography

Daniel F. Macedo is a Ph.D student at Université Pierre et Marie Curie-Paris VI. He received a M.Sc. and a B.Sc. degrees in Computer Science from Federal University of Minas Gerais (UFMG), Brazil, in 2004 and 2006, respectively. Daniel participated in the program committee of IEEE ISCC 2006 and 2007 and ICWMC 2006 and 2007. He received the “Best Simulation Article” award at IEEE ISCC’05 for the paper “A Pro- Active Routing Protocol for Continuous Data Dissemination Wireless Sensor Networks”. His main research interests are wireless

networks, distributed algorithms, network management, and peer-to-peer computing.

Aldri L. dos Santos is a researcher and professor at the Federal University of Paraná, in the Department of Informatics. He received his Ph.D. in Computer Science from Department of Computer Science of the Federal University of Minas Gerais, Belo Horizonte, Brazil. Aldri received both his M.Sc. and B.Sc in Informatics from Federal University of Paraná, Curitiba, Brazil. From July to October 2001 he was an invited researcher at Cyber Solutions Inc., Sendai, Japan. He also is member of the SBC (Brazilian Computing Society).

Guy Pujolle received the Ph.D. and "Thèse d'Etat" degrees in Computer Science from the University of Paris IX and Paris XI on 1975 and 1978 respectively. He is currently a Professor at the Pierre et Marie Curie University (Paris 6) and a member of the Scientific Advisory Board of Orange/France Telecom Group. He spent the period 1994-2000 as Professor and Head of the computer science department of Versailles University. He was also Professor and Head of the MASI Laboratory (Pierre et Marie Curie University), 1981-1993, Professor at ENST (Ecole Nationale Supérieure des Télécommunications), 1979-1981, and member of the scientific staff of INRIA, 1974-1979. He is currently an editor for International Journal of Network Management, WINET, Annals of Telecommunications, and IEEE Surveys & Tutorials.