Context-Aware Information Delivery in Assistive Environments over a Publish-Subscribe Internet

George C. Polyzos
Mobile Multimedia Laboratory, Department of Informatics
Athens University of Economics and Business
11362 Athens, Greece
e-mail: polyzos@aueb.gr, http://mm.aueb.gr/

***Abstract*—Future Internet technologies are offering new capabilities for service and data delivery. We present the potential for enabling the context-aware content adaptation and specialized delivery of healthcare information of the Publish-Subscribe Internetworking (PSI) Future Internet architecture. The Information-Centric Networking paradigm of PSI brings information at the center of the approach, rather than muddying it with the 'who' and 'where' of the information, or data. In addition, it enables seamless information morphing, effective access control policies, and other security features.**

Keywords; Future Internet; Information-Centric Networking; content adaptation; healthcare; security

I. INTRODUCTION

The introduction of the pervasive healthcare paradigm has focused attention, in addition to chronic patients under constant medical supervision, also towards the elderly and the general population, in particular habitants at remote, isolated and underserved locations. In this context, advanced electronic healthcare devices and services that were once available only at big hospitals, are now expected to be available through a communications network anytime, anyplace, and to anyone. A medical assistive environment concerns the utilization of pervasive and ubiquitous technologies for delivering the above services. The word "assistive" may also be used in a more general setting, including not only assisting persons with recognized health problems, but also as empowering any human to improve quality of life by enhanced sensing and computer-based health support.

For example, the ARCHANGEL project is about a sensor-based system that focuses on monitoring the health of the elderly and people with special needs. The system relies on location sensors and GPS-based location tracking. Its main objective is to detect changes in the health status of the individuals that it monitors and to provide alerts and a preliminary diagnosis as quickly as possible whenever something out of the ordinary occurs. This applies both inside the home using sensors and outside the home using positioning-enabled cellular phones of the monitored individuals. A secondary function of the system is the actuator-based automation of certain tasks inside an individual's home, which makes use of the history of an his actions to automatically build a profile of the individual's "standard" behavior. Alerts are propagated over the cellular network and the Internet to appropriate servers, where relevant filtering software can decide whether or not to forward the alert to qualified human caretakers [1]. The privacy and security implications of the above scenario are obvious and substantial.

We see that the type of information that is exchanged in assistive environments is characterized by high diversity, the need for adaptation (morphing) and stringent security and performance requirements. In such an environment, the notion of Future Internet (FI) and the new concepts and architectures that it brings with it, such as distributed storage and processing and various flexible networking technologies, such as ad hoc, delay/disruption tolerant etc. as well as cloud and pervasive computing and exploitation of the Internet of Things can fit perfectly, enabling the development of advanced assistive applications.

In this paper we discuss the trends in FI technologies and architectures enabling assistive environments focusing on advanced networking and context awareness. We particularly focus on a promising *clean-slate* information oriented approacg, the *Publish-Subscribe Internet* (PSI) architecture. In PSI 'everything is information and information is everything,' i.e. information is the main building block for realizing its goals and recursion is exploited. PSI is based on the so-called publish-subscribe paradigm, i.e., its communication model is based on the *publish* and *subscribe* primitives rather than the traditional *send* and *receive* primitives. Information-centric architectures appear to be ideal for pervasive healthcare environments as they achieve effective information collection, dissemination, processing and governance [2].

II. LIMITATIONS OF THE CURRENT INTERNET

The problems of the current Internet come as a natural consequence of its architecture being designed to address the communication needs back at a time when a network was needed for the sole purpose of sharing rare, expensive resources such as cycles of expensive mainframes. The basic requirement from the Internet infrastructure at that time was merely that of forwarding packets of data among a limited number of stationary machines with established trust relationships. The key design principles of the Internet were such that it has been very simple to link new networks to the Internet and have allowed a tremendous growth of its size. In parallel to the Internet's growth, an unprecedented number of innovations in the applications and services running on top of it and in technologies below the networking layer have emerged. This is attributed to the fact that Internet's protocol architecture followed the hourglass approach, where the networking layer forms the waist of the hourglass and is powerful enough such that almost any application can run on top of it and simple enough such that it can run over any link-layer technology.

These merits of the Internet's architecture have facilitated the tremendous size growth of Internet as well as the introduction of new applications to fulfill emerging needs. However the Internet was never designed to address the new requirements and in order to help it "evolve," a vicious cycle of functionality patches began.

ACWR '11, Dec 18-21 2011, Amritapuri, Kollam, Kerala, India
ACM 978-1-4503-1011-6/11/12.

Most of those patches proved to be only partial and temporal solutions and many current but also emerging requirements cannot be addressed by the current Internet architecture. This has raised the issue of whether a new, *clean-slate* architectural approach for the Internet is actually needed. Along this lines a research community has been formed which has clarified the limitations of the current Internet, is discussing the key requirements and objectives of the Future Internet, and started proposing new architectures to address them.

Next we will briefly discuss some key problems and limitations of the current Internet architecture and technology particularly relevant for wireless networks, e-health applications and assistive environments and then outline the other key problems.

A. *Problems in Supporting Mobility, Wireless Terminals and Devices with Limited Resources*

The addressing philosophy (and basic protocols) of the Internet was designed only with stationary (fixed) hosts in mind, binding it to location and network topology. However, current internet statistics show a constantly increasing number of non-fixed hosts accessing the Internet and forecasts reveal that by 2015 traffic from wireless terminals will exceed traffic from wired terminals [3]. Wireless and mobile devices may easily switch networks, changing their IP address and thus introducing new networking modes based on intermittent and possibly opportunistic connectivity. The Mobile IP protocol which came as a patch to remedy the problem of locating moving hosts implies a "triangular routing model": packets need to be routed first to the home agent of the mobile host at its respective home subnet, and from there to the current location of the mobile node. This is the source of a major inefficiency since traffic has to travel a longer path than the optimal one and the problem is aggravated especially when the mobile node, its home agent and the third party (data source or destination) are located in distant Autonomous Systems (AS) in terms of network proximity. Although the IPv6 version of the protocol allows the direct communication of a mobile node with remote third parties, the "triangular routing model" is still present in the control plane, resulting in high delay for handoffs. Moreover, Mobile IP employs tunneling thus incurring high processing load at the routers.

In addition, there is a known problem regarding accessing, through Mobile IP, services based on ingress filtering, i.e., services that demand that incoming traffic comes from the actual network it claims to originate from. Due to Mobile IP these services see the home address of the mobile node when sending packets which is different than the care-of address assigned by the guest network when receiving traffic back from the mobile node. Furthermore, the Mobile IP solution does not comply with ISP inter-domain policies. Just like any overlay network, traffic tends to violate Border Gateway Protocol (BGP) rules because it is routed first to the home agent of the mobile node and from there it is re-routed to its currently hosting network. This is responsible for (a) "valley routing," i.e., a client AS serving traffic from a provider AS, or traffic originating from some peer AS's related AS, and (b) "exit policies violation," i.e., traffic exiting from a different exit point of the source network than the exit point it is supposed to according to the BGP rules for a given traffic destination.

Additionally, because TCP cannot distinguish between packet losses due transmission errors much more prevalent over wireless links and packet drops due to congestion and reacts badly, it performance can be severely degraded in wireless and even more in mobile environments (because of additional problems) [4].

Finally, the TCP/IP stack is considered heavy for many devices with limited resources, typically found in sensor or nano-scale networks, where processing, storage, and transmission capacity can be severely constrained [5].

B. *Problems in Security and Trust*

The Internet protocols and the overall architecture were initially designed for operating in a completely trusting and cooperative environment. User and data authentication, integrity and privacy were not a requirement and the focus was on openness and flexibility in allowing new hosts to join the network. Moreover, the protocols and algorithms were designed to forward all traffic injected in the network, resulting in an imbalance of power between senders and receivers. These characteristics have allowed spammers, hackers and attackers in general to relatively easily perform Denial of Service attacks against the Internet infrastructure, or against Internet hosts and services, or to obtain private data while easily covering their tracks. In order to cope with such malicious behaviors add-on security patches and trust mechanisms have been developed. Internet techniques such as firewalls, NATs, and spam filters along with security protocols have been introduced. However, such solutions do not penetrate deep into the network and bad data still get forwarded, clogging systems and possibly fooling the filtering technology. The additional processing requirements and the Internet's end-to-end philosophy have so far blocked the placement of security and trust mechanisms deeper into the network, where it would be most efficient and effective to identify and stop such attacks. Many of these problems are largely due to the disconnection between information semantics at the application layer and opaque data in individual (IP) packets. This disconnection places a significant burden on integrating accountability mechanisms into the overall architecture. Point solutions such as deep packet inspection or lawful interception try to restore this broken link between the actual information (semantics) and the scattered data in individual packets. However, this is achieved at a relatively high cost and is therefore only applied for imminent or important problems, usually by or on the request of law enforcement. Thus we have reached a point where we may have secure protocols, but the overall Internet is still not adequately protected against malicious attacks. At the same time the lack of an accountability framework which would allow non-intrusive and non-discriminatory means to detect misbehavior and mitigate its effects while keeping open and broad accessibility to the Internet and ensuring communication privacy (hiding from non-authorized parties even that communication took place) is a crucial limitation to overcome [5].

C. *Additional Problems*

A number of other important problems with the current Internet have also been identified; outlined, these are:

- routing scalability problems
- problems in congestion control
- problems with content distribution
- problems in providing Quality of Service
- management and control problems

III. The Quest for Future Internet Technologies

The need for designing a new architecture for the Future Internet has led the research community towards various directions and, as a result a plethora of different research projects in both Europe, US and ASIA have emerged [6].

IV. Information-Centric Networking

A. *Key Concepts & Principles of ICN*

One of the most promising directions, as revealed by the fact that many promising FI projects have focused on it, is the Information-Centric Networking (ICN) paradigm. The popularity of ICN is attributed to the fact that the Internet's role has radically

changed, from an infrastructure for sharing distributed resources, to an infrastructure for (mostly) delivering content or information. In [7] the authors identify a list of basic requirements set by content-centrism that the current internet architecture cannot adequately satisfy:

- Name persistence. Mechanisms such as HTTP redirect and dynamic DNS are used for this purpose, but are not sufficient.
- Authenticity. Current mechanisms focus on securing the communication channel rather than authenticating the data itself.
- Availability.

The first ideas for changing Internet's design nature from host-centric to content-centric was introduced almost a decade ago in seminal papers by Carzaniga & Wolf, e.g. [8], on content-based networking. This approach decouples the data from their sources or current location(s) by means of the location-identity split in naming. The basic assumption behind this is that the content is identified, addressed, and matched independently of its location anywhere in the network [9]. In ICN, instead of specifying the source-destination pair that communicates, the piece of data is identified.

An indirect implication (and also benefit) from moving from the host identification regime to the information identification regime is also that information (data) retrieval becomes now receiver driven. In contrast to the current Internet where senders have the absolute control, in ICN no data can be received unless it is explicitly requested by the receiver. After a request is sent, the network is responsible for locating the desired data, by routing the request to the best location where the data is available using anycasting mechanisms.

Since ICN inherently and effectively supports caching in network elements (in-network caching) the network may satisfy a data request not only through locating the actual source, but also by involving in-network caches that hold copies of the requested data (or pieces of it). Caching is considered a core service of the network. This is similar to embedding the functionality provided by CDNs inside the network.

Caching can refer to caching information objects as a whole, or to packet-level caching. Depending on the proposed architecture, caching can be seen as an enhancement that makes data exchange more efficient (faster, less costly), or as the main service of the network as in the case of CCN [10]. No matter the case, ICN-based architectures see non-opaque data transfers in the sense that flows are identified based on the information they carry. Therefore, information fragments (packets in current terms) can be cached and retrieved easily unlike in the case of deep packet inspection with IP, which is costly and impossible with encryption (or potentially inaccurate if based on timing and other inferences). Additionally, access control to data can be applied directly at the network layer with ICN, limiting the propagation of content.

B. ICN Realization through the Publish-Subscribe Model

Many ICN architectures are based on the publish-subscribe (pub/sub) communication model. In pub/sub, senders "publish" content (at the network level, i.e., they advertise the availability of content). Interested receivers then need to "subscribe" to the publication, i.e., express their desire for a specific publication/content (whenever available). Subscriptions may precede publications or the other way around. Inside the network, brokers are responsible for matching subscriptions with publications i.e., provide the rendezvous function. No one receives any content for which they have not explicitly requested by means of a subscription (except for the rendezvous network, at the signaling level, i.e., the brokers that receive publications and subscriptions). Data forwarding follows on the instruction of the rendezvous point (network) to the publisher (one, of the potentially many, chosen based on various criteria after consulting an associated topology manager). The process is illustrated in [11]. Note that forwarding can be achieved in a way that can ensure that the publisher does not know the end subscribers (and with easy multicast support), and with dynamic paths so that the publisher cannot in the future hope to be able to reach the subscribers. A method to achieve this is described in [12]. As a by-product this avoids or limits (D)DoS attacks (except to the rendezvous network [13] that needs special attention).

The pub/sub model gains momentum due to its inherent advantages that include (a) information-centrism, i.e., better match to prevalent and important applications (b) decoupling in space and time with respect to information sources/producers and sinks/consumers, (c) support for mobility as discussed in [14], [15] and [16], and (d) support for anonymity [17].

The information-centric character of the pub/sub model stems from the focus on the rendezvous (or resolution) function. The matching of publications with subscriptions is primarily based on the identity of information objects, thus centering communication around information.

The strength and also the largest benefit of the pub/sub communication model stems from the fact that publication and subscription operations are decoupled in time and space. The communication between a publisher and a subscriber does not need to be synchronized, i.e., the subscriber may publish events when the subscriber(s) are offline and the subscriber(s) may get notified about events while the publisher of the event is disconnected. The publishers do not usually hold references to the subscribers, neither do they know how many of these subscribers are participating in the interaction and similarly, subscribers do not usually hold references to the publishers, neither do they know how many of these publishers are participating in the interaction (space decoupling). The fact that the publication-subscription matching takes place at an independent point in the network along with the time/space-decoupling, reduces synchronization requirements between the participating entities and thus allows for efficient support of multicast, mobility, as well as multihoming and indirection [13].

The goal of many current ICN efforts is to result in clean-slate, natively supported architectures. However, overlay and mixed modes co-existing or over the current Internet are also considered, particularly as effective deployment strategies; e.g., see the approach in [18].

C. The Role of Satellite Communications

Key benefits and limitations of SatCom include globally applicable traits, such as the inherent broadcasting and multicasting capability, very wide area coverage, ubiquitous services etc., which are uniformly beneficial to all FI technological solutions.

Any FI initiative should effectively integrate all existing networking infrastructure in a common framework. Moreover, FI initiatives may actually help the role of satellite networks grow in presence and impact in tomorrow's Internet by facilitating SatCom integration or coexistence with terrestrial networks. In this context, we examine in an ESA study how SatCom can integrate with other networks in the context of the FI [19].

Efficient support for mobility has been identified as one of the main drivers for FI initiatives. In this context, efficient support of mobile users is of high significance for evaluating the integration suitability of FI techniques with SatCom networks, which provide a natural framework for supporting user mobility. Several proposed FI techniques suffer from limitations in various aspects of host mobility, most notably during mobility of a publisher. Satellite networks can provide a framework for mitigating these problems due to their wide coverage areas and centralized architecture.

V. ICN AS AN ENABLING TECHNOLOGY FOR E-HEALTH AND ASSISTIVE ENVIRONMENTS

Context-aware assistive environments today face difficulties due to the limitations of the current Internet design. Such systems demand complex information manipulation and effective security mechanisms which currently have to be provided with add-on solutions. This is necessary as users utilize various devices for creating and retrieving information and because the generated content contains sensitive data which has to be properly secured.

In [20] we presented a conceptual solution for a context-aware assistive environment based on the PSI architecture. We showed that by designing an architecture around information items and by providing functions for manipulating and securing them, pervasive health assistive solutions can be more easily deployed. Moreover we implemented part of this conceptual architecture utilizing a PSI prototype. Although at a small scale, the implementation gives a hint about the scalability, extensibility and security features of such an architecture.

As the PSI prototype evolves, it is expected that functionality currently implement by add-on solutions, will be part of the architecture's core functionality. In the current prototype the rendezvous function is performed by a single machine. Future releases of the prototype will include distributed rendezvous functionality enabling more secure and scalable handling of publications and subscriptions. Moreover, currently the topology manager of the prototype creates paths among end-points without taking into consideration user preferences. It is expected that in future releases it will be possible for a publisher or a subscriber to include in their requests, e.g., nodes through which data should be sent. This new feature will enable the deployment of content transformation nodes—subscribers or publishers wanting content transformation will require the delivery paths to include these nodes.

In the developed solution public keys were used to define access control policies. Future work in this domain includes the possibility of definition of access control policies using more attributes. Such attributes can include for example location and policies such as "in case of emergency, all nearby hospitals must be notified" should be possible to be specified and realized by the system.

The API created for this implementation allows the creation of a single scope per patient. Future versions will allow the creation of a hierarchy of scopes allowing for better organization of information. Finally, patient monitor devices are at present communicating directly with the storage devices owned by the patient using HTTP. It is in our future plans to allow device interaction with the core PSI network. If patient devices are capable of communicating with the PSI network, then the storage devices can be regarded as a network service—such as cloud storage, or caches.

ACKNOWLEDGMENT

The research reported here has been supported in part by FP7 project PSIRP, under contract ICT-2007-216173, by FP7 project PURSUIT, under contract ICT-2010-257217, and by project ARCHANGEL, through a *Cell Phone as a Platform for Healthcare* award from Microsoft Research.

REFERENCES

[1] G.J. Papamatthaiakis, G. Xylomenos, and G.C. Polyzos "Monitoring and Modeling Simple Everyday Activities of the Elderly at Home," Proc. 7th Annual IEEE Consumer Communications and Networking Conference (CCNC 2010, special session on 'Advanced Technologies for Care at Home'), Las Vegas, NV, Jan. 2010.

[2] D. Trossen, D. Pavel, K. Guild, J. Bacon, J. Singh, "Information-centric Pervasive Healthcare Platforms," PervasiveHealth, Mar. 2010.

[3] Cisco Systems, "Cisco Visual Networking Index: recast and Methodology, 2010–2015," white paper, June 2011.

[4] G. Xylomenos, G.C. Polyzos, P. Mähönen, and M. Saaranen, "TCP Performance Issues over Wireless Links," *IEEE Communications Magazine*, vol. 39, no. 4, April 2001.

[5] Future Internet Architecture (FIArch) Experts Group, European Community, "Fundamental Limitations of current Internet and the path to Future Internet", white paper, Mar. 2011.

[6] P. Mähönen, D. Trossen, D. Papadimitriou, G.C. Polyzos, D. Kennedy, eds., "The Future Networked Society: a White Paper from the EIFFEL Think-Tank," Dec. 2006. http://www.fp7-eiffel.eu/fileadmin/docs/EIFFEL-FINAL.pdf

[7] T. Koponen, M. Chawla, Byung-Gon Chun, A. Ermolinskiy, Kye Hyun Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," Proc. ACM SIGCOMM Aug. 2007.

[8] A. Carzaniga and A.L. Wolf, "Forwarding in a Content-based Network," Proc. ACM SIGCOMM, Karlsruhe, Germany, Aug. 2003.

[9] S. Arianfar, P. Nikander, and J. Ott, "On Content-centric Router Design and Implications," Proc. ACM ReArch Nov. 2010.

[10] M. Diallo, S. Fdida, V. Sourlas, P. Flegkas, L. Tassiulas, "Leveraging Caching for Internet-scale Content-based Publish/Subscribe Networks," Proc. IEEE ICC, Kyoto, Japan, June 2011.

[11] N. Fotiou, G.C. Polyzos, D. Trossen, "Illustrating a Publish-Subscribe Internet Architecture," *Telecommunication Systems*, Springer, vol. 52, no. 3, Special Issue on 'Future Internet Services and Architectures: Trends and Visions,' Mar. 2013 (online on 23/2/2011: http://www.springerlink.com/content/t6m0k022042088t5/fulltext.pdf)

[12] P. Jokela, A. Zahemszky, C.E. Rothenberg, S. Arianfar, P. Nikander, "LIPSIN: Line Speed Publish/Subscribe Inter-networking," Proc. ACM SIGCOMM, Barcelona, Spain, Aug. 2009.

[13] N. Fotiou, G.F. Marias, G.C. Polyzos, "Towards a Secure Rendezvous Network for Future Publish/Subscribe Architectures," 3rd Future Internet Symposium, Berlin, Germany, Sept. 2010.

[14] N. Fotiou, K. Katsaros, G.C. Polyzos, M. Sarela, D. Trossen, G. Xylomenos, "Handling Mobility in Future Publish-Subscribe Information-Centric Networks," *Telecommunication Systems*, Springer, Special Issue on 'Mobility Management in the Future Internet' (to appear).

[15] V.A. Siris, X. Vasilakos, and G.C. Polyzos, "A Selective Neighbor Caching Approach for Supporting Mobility in Publish/Subscribe Networks," Proc. 5th ERCIM Workshop on eMobility, Catalonia, Spain, June 2011.

[16] V. Giannaki, X. Vasilakos, Ch. Stais, G.C. Polyzos, and G. Xylomenos, "Supporting Mobility in a Publish Subscribe Internetwork Architecture," Proc. ISCC, Kerkyra, Greece, June 2011.

[17] N. Fotiou, G.F. Marias, and G.C. Polyzos, "Publish-Subscribe Internetworking Security Aspects," in *Trustworthy Internet*, N. Blefari-Melazzi, G. Bianchi, L. Salgarelli, eds., Springer, May 2011.

[18] K.V. Katsaros, G. Xylomenos, G.C. Polyzos, "MultiCache: an Overlay Architecture for Information-Centric Networking," *Computer Networks*, vol. 55, no. 4, pp. 936-947, Elsevier, Special Issue on 'Architectures and Protocols for the Future Internet,' T. Wolf & L. Eggert, eds., Mar. 2011.

[19] K. Liolis, V.A. Siris, and G.C. Polyzos, "On Satellite-assisted Publish-Subscribe Future Network Architectures for Smart M2M Applications," 6th Future Internet Cluster Workshop, 'Network Architecture for a Smarter Environment,' Poznan, Poland, Oct. 2011.

[20] Ch. Doukas, N. Fotiou, G.C. Polyzos, I. Maglogiannis, "Context-Aware Delivery of Information in Assistive Environments utilizing Future Internet Technologies," Proc. 4th ACM International Conference on Pervasive Technologies Related to Assistive Environments (PETRA), Crete, Greece, May 2011.