



A Content-Centric, Publish-Subscribe Architecture delivering Mobile Context-Aware Health Services

Panagiotis K. GKONIS¹, Charalampos Z. PATRIKAKIS², Angelos-Christos G. ANADIOTIS², Dimitra I. KAKLAMANI², Maria Teresa ANDRADE³, Andrea DETTI⁴, Giuseppe TROPEA⁴, Nicola BLEFARI MELAZZI⁴

¹*Singular Logic*

Email: pgonis@singularlogic.eu

²*National Technical University of Athens, Department of Electrical and Computer Engineering, Emails: aca@icbnet.ntua.gr, bpatr@telecom.ntua.gr, dkaklam@mail.ntua.gr*

³*INstituto de Engenharia de Sistemas e Computadores (INESC) do Porto, E-mail: mandrade@inescporto.pt*

⁴*Consorzio Nazionale Interuniversitario per le Telecomunicazioni
Emails: andrea.detti@uniroma2.it, giuseppe.tropea@cnit.it, blefari@uniroma2.it*

Abstract: The goal of this paper is to report on the design of a novel content-centric future Internet platform that can support added value services incorporating mobility, context awareness and enhanced security and privacy. We present the proposed platform together with a use case study, dealing with the provision of added-value health services, consisting in the safe and accurate management of medicine prescriptions. Both the design of the platform and the use case are the results of work performed in the context of the European Research project Convergence [1].

Keywords: Future Internet, Context Centric Networks, Context-aware services, e-health.

1. Introduction

It is expected that future Internet applications will support interoperability of heterogeneous mobile technologies as well as context-awareness. These aspects have been addressed in the literature with a wide set of proposed solutions based on the plain TCP/IP API [9][10]. Nevertheless, there is an increasing interest in replacing or upgrading the actual host-centric and request-response TCP/IP API, with a new API that is content-centric [5] and publish-subscribe [8]. With this regard, this paper deals with the design of a mobile context-aware application on top of a publish-subscribe content-centric system.

By using a content-centric API, a user can retrieve remote content by feeding the interface with a description of the content (e.g., its name), rather than providing the location (i.e. a URL) of the host that has the content. By exploiting the publish-subscribe interaction model, a user can express interest in a certain content and get notified when such content or a succession of updates will become available on the network, while she does not obtain an error from the API if content is not available at the instant of her request.

The publish-subscribe content-centric system architecture we consider here is the one we are developing in the CONVERGENCE project [1]. A classic content-centric system focuses on providing access to contents (documents, media, files etc); however, we argue that a future Internet architecture should also provide effective means to access interactive

services (i.e., HTTP sessions, Telnet, voice calls, etc.). For this reason, with respect to other content-centric solutions, in CONVERGENCE we deal not only with contents but also with service access points, in the sense that both content and access points are identified by a name and referred to with such name, instead that with their locations. In this sense, CONVERGENCE may be seen as an extension of literature content-centric solutions. This implies that we must generalize the concept of content and introduce a new term that comprises both content and access points. Such new term is “resource” and by resource we mean both content and access points. Resources are described by a set of metadata contained in a common structure, which we name Versatile Digital Item (VDI). The VDI defines a package of digital information with a unique identifier, independent of the machine where the VDI is hosted. VDIs will be designed to handle all possible kinds of resources, independent of their structure or geographical location. Resources will be made available through the creation and publication of the corresponding VDIs. Users will be able to access resources by subscribing to them, either by describing their interests to the system via semantic queries that characterize the resources or by explicitly requesting a given resource by means of its name. Published VDIs that match these subscriptions will be delivered or made available to users. This architecture can be used to enhance applications and support real time services.

The CONVERGENCE system is exploited to deliver a mobile context-aware application, which includes fast and accurate medical purchase and prescription. The application platform represents medicines as separate VDIs, which include all relevant information such as different names of the medicine per country of export, indicative price, possible side effects and availability. Doctors and patients will be able to access this information with the use of their mobile device anywhere in the world by simply subscribing to the CONVERGENCE system.

The rest of this paper is organized as follows: Section 2 introduces the VDI concept, highlighting its major features and provides a description of the Convergence architecture. In Section 3 the proposed context-aware platform for mobile medical applications is introduced. Implementation possibilities are highlighted in Section 4. Finally, concluding remarks are made in Section 5.

2. CONVERGENCE Fundamental Aspects

2.1 The Versatile Digital Item concept

The Versatile Digital Item (VDI) is a structured digital object, which allows defining a virtual container for all data or components associated to a given resource. Resources can be natively digital, or can have a digital representation when they are real-world objects or entities, such as persons, animals, or organizations. Moreover, if the resource is not a service access point, it might also be fully embedded inside the VDI that describes such resource

The VDI concept is based on the MPEG-21 Digital Item (DI) specification. A DI can be seen at the conceptual level as a package that bounds together multimedia resources with associated descriptions (e.g. rights descriptions, semantic descriptions, low-level audio and video descriptors, etc). Part 2 of the MPEG-21 standard [1] specifies a method based on XML Schema for building a document, the Digital Item Declaration (DID), containing the declaration of the structure of the DI. Inside this XML document (the DID), standard MPEG-21 mechanisms are used to either convey resources and descriptions embedded directly in the DID or to provide information on the location of the resources to be fetched. The DID provides an indication of the composition of the DI in terms of media resources and metadata, as well as of the relationships among its several constituents. The DID

typically conveys metadata only. The VDI extends the DI concept by enabling the declaration of any type of resources, including services, real-world objects and people, and by supporting a whole new set of semantic and manipulation operations on the VDIs.

Every VDI is uniquely identified by a given VDI identifier, assigned to the specific VDI when the VDI is first created and finalized as a self-contained piece of information, to be injected in the CONVERGENCE system and traded between communicating peers.

Each VDI may also include a License. Among other things, the license to a VDI defines its access rights, stating its owner, its authors, who can modify it. The license could state, for example, that only the authenticated author can modify the VDI, that other authors can add resources to the VDI without deleting the original resources or that other authors can freely replace resources in the VDI. The act of modifying a VDI is in fact the creation of a new VDI, semantically related with the original, but with a new digital signature recalculated on the whole updated package of bits, as well as with a new identifier. This new VDI can thus be trusted, based on the author of the change, and guarantees that the original VDI has not been tampered with after the publication.

Convergence wants to meet the goal of designing a VDI which is a trusted self-contained package of information that travels the network and is composed of a dynamic part which can be manipulated as a shared resource, where actors can add or replace new pieces of information, and a semantic part that contains (among other things such as the creator) the metadata and a stable identification block which links back, through a level of indirection, to this specific VDI regardless of the successive manipulations. These VDI updating operations are designed in a way that implies no alterations to already published VDIs, but allows for removal of obsolete VDIs (enabling *digital forgetting*) and for retrieval of all old versions of a VDI, if requested by the user.

2.2 The CONVERGENCE System Architecture

The CONVERGENCE system architecture is structured into three different Functional Levels as depicted in the schematic diagram of Figure 1. Each of the Convergence levels, namely Application, Content and Network, aggregates a distinct set of conceptually similar functions, which in turn are implemented by one or more Functional Blocks. Four Functional Blocks have been defined: the CONVERGENCE Network (CoNet), the CONVERGENCE Middleware (CoMid), the Community Dictionary Service (CDS) and CONVERGENCE Applications. A CONVERGENCE system is formed by a set of interconnected devices implementing one or more Functional Blocks. CONVERGENCE devices can be *peers* or *nodes*, depending on the Level of the Convergence stack deployed in them. A *peer* deploys both the Content Level functionality and the Network Level functionality, whereas a *node* deploys Network functionality only.

We argue that implementing all the functionalities of a content-centric publish-subscribe system at network level overloads processing capabilities of backbone nodes of such a future Internet. Therefore, we follow the traditional approach to put complexity on the edge. When VDIs and resources are handled by the CONVERGENCE Network (CoNet) they become CoNet *named-resources*. We place at network level only simple functionalities that allow network nodes to identify and locate them in terms of their “name”; i.e. the CONVERGENCE network level follows a *networking named data* paradigm [3]. Thus handling of resources on the basis of “what” they contain (rather than on the basis of their names) is a complex function that exploits VDI metadata and that we place at *content level* (CoMid), i.e. on CONVERGENCE peers. Furthermore, the content-level exports a publish-subscribe API to the Application level.

On the other hand, the network-level provides access to CoNet named-resources on a public or private network infrastructure [3][4][5]. A CoNet named-resource is anything that

can be identified by means of a name. Examples of CoNet named-resources include: a VDI, an electronic document, an image, a source of information with a consistent purpose, the point of access to a service, and a collection of other resources. The name of a named-resource is its *network-identifier* (NID). Unlike Internet URLs, which include information about “where” an object is located, CONVERGENCE NIDs do not contain any reference to location. A network-identifier is an *anycast* address: a system may contain multiple replicas of the same named-resource. Network functionality will connect a user to the “best” replica (e.g., the closest one). For example: a PDF copy of the TIMES newspaper for Sept-30-2010 could be a CoNet named-resource identified by “the times: Sept-30-2010”; the service-access-point for an SQL database service provided by the Foo company could be a CoNet named-resource identified by the name “Foo: sql-database”. The handling of CoNet named-resources through their network identifiers, rather than IP addresses, means that the network infrastructure is aware of the name of the resources that it is handling. This awareness can be exploited to support anycast routing, replication, and in-network caching [6] (if resources are data and not points of access to a service). These mechanisms are very needed but not supported by current IP networks. The possibility of accessing CoNet named-resources without having to specify their location, simplifies mobility issues. It also means that their names can be portable and do not depend on the service provider handling the CoNet named-resource.

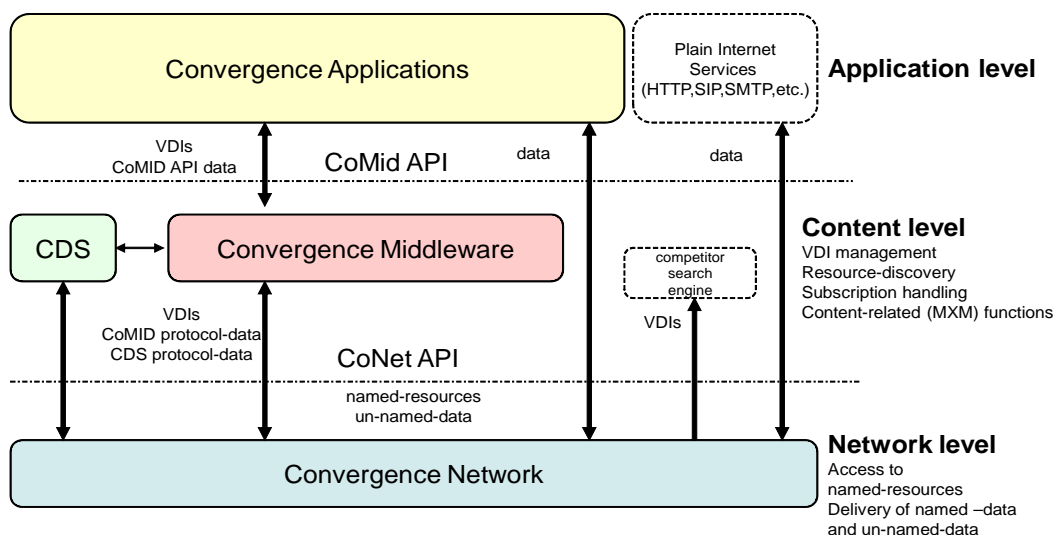


Figure 1: The Convergence Network Architecture (from deliverable 3.1 of CONVERGENCE project)

The CoNet allows users (e.g., end-users and providers) to advertise and revoke their named-resources. A CoNet named-resource can be replicated in different geographical locations by adopting the same network-identifier. Replication enables users to retrieve/exploit CoNet named-resources that are closer to their location and improves reliability of access. The CoNet provides *secure distribution* of CoNet named-resources by assuring that the network-identifier cannot be forged and that receivers can verify the validity and provenance of CoNet named-resources [4][7]. Network functionality may autonomously cache transiting CoNet named-resources [6], if this is allowed by security rules and by the nature of the resource.

When a CoNet named-resource consists of actual data, the CoNet provides the means to deliver it to intended recipients. It can also facilitate the delivery of the data through replication/caching. When a CoNet named-resource is a service-access-point, the CoNet provides the means to exchange information (*un-named-data*) between a requesting upper layer entity and the upper layer entity addressed by such named-service-access-point. This functionality is needed by “traditional” client/server services (e.g. HTTP, POP, SMTP).

CoNet functions are supported by a novel network protocol, named CoNet Protocol (CP). Very briefly, (more details to be found in CONVERGENCE deliverables, [1]) CP is located on top of a layer-2 (e.g. Ethernet, MPLS) or a layer-3 protocol (e.g., IPv4 or IPv6). CP is connectionless, uses packets that include the NID and security information, and carries out name-based routing, caching and encapsulation functionality. CP packets transport: i) requests of named-resources, ii) un-named-data and iii) named-data. In the latter case, a CP packet contains a chunk (e.g., 512 kB) of the whole content (unless the whole content fits in a chunk) and CoNet nodes may cache these packets to deliver them to requesting entities/users (supporting in-network caching). CoNet nodes support name-based routing, in which CP packets are forwarded on the basis of the NID. The routing functions aim at discovering the “underlying” address (e.g. IP or Ethernet) of the device toward which the CP packet has to be forwarded. Such discovery can be done either following a route-by-name approach [4] or a lookup-by-name approach [4]. Encapsulation functions insert the CP packet in the underlying data unit and, if needed, carry out segmentation and reassembly functions.

In case the underlying technology is IPv4 (or IPv6), a part of the CP packet header (e.g., the NID) is transported by a novel IPv4 option (or by an IP v6 extension header); CoNet nodes are enhanced IP routers able to parse and interpret such option or extension header.

To support fast route-by-name [4], we propose that the IP Forwarding Information Base (FIB) can include also the NIDs of named-resources. Thus the FIB is an *hybrid* one, in the sense that contains both traditional IP network addresses and NIDs. The NID entries in the FIB point to the address of the next device in the path towards the node where the content is located. Since NIDs can be very numerous, the set of NIDs reported in the FIB can be limited e.g. by including only the most popular NIDs.

Finally we propose that the FIB contains also the NIDs of contents stored in the local cache, to support fast cache lookup. In this case, the FIB entries point to the local cache. In this way, CP/IP packets can be processed by Hardware Engines (fast-line) and avoid that the router CPU has to process all CP/IP packets.

Intermediate functionalities between CoNet and CONVERGENCE-enabled Applications are provided by the Content Level, through the CoMid and the CDS Functional Blocks. This level is responsible for handling resources based on their content. In CONVERGENCE, as referred above, resources are represented as VDIs, hence the Content Level essentially deals with metadata, enabling the generation of metadata associated to resources, or encapsulation of existing metadata, inside the VDIs and the interpretation of such metadata to facilitate finding resources that match users’ subscriptions or requests.

Accordingly, CoMid delivers functionality to “understand” and manipulate VDIs (create, browse, publish, search, etc...) and the CDS functionality to extract knowledge from the metadata contained in or pointed to by VDIs.

The CDS can be regarded as the CONVERGENCE knowledge base. The CDS will contain the ontologies used to semantically describe any resource, from publishing a VDI with this resource to subscribing to VDIs by performing semantic search. When the user subscribes to a set of VDIs satisfying certain search criteria, the CDS will provide CoMid with all matching concepts, coming from domain or user-defined ontologies. CoMid will then perform matching with the VDIs and deliver the desired information to the user.

Finally, the Application Level of CONVERGENCE comprehends a number of pre-packaged Tools to provide a friendly environment for the user to interact with VDIs, making use of the API exposed by the Content Level. It offers a variety of high-level functionalities to users, such as fast and easy VDI retrieval and modification as well as context-aware applications, hiding the technical details associated to VDIs.

This architecture has been devised to provide a simple and user-friendly, yet powerful, environment to consume any type of resource and make it available based on semantic descriptions. Modifying or updating published resources is achieved by creating the required links to the previous versions, through the combined use of semantic descriptions at the CoMid and of network-identifiers at the CoNet. When submitting a resource, a corresponding VDI is created and then published into the system. During the creation process, metadata will have to be added, digital rights associated to the resource, management of the Intellectual Property included, etc. To accomplish this, the Application Level uses, on behalf of the user, the API offered by the CoMid. Given the semantic-based publish-subscribe approach of CONVERGENCE, any VDI publication must include a semantic analysis phase, so that the VDI is efficiently stored and replicated in the system, according to its semantic attributes. This is the task of the CDS. This semantic analysis is also an essential step in order to find already published resources. Likewise, subscriptions are typically made using rich semantics and therefore also require the intervention of the CDS. Although possible, typical subscriptions are not made for a specific resource but rather for resources of a given type, characterized by their semantic descriptions.

3. Use Case: Fast and Accurate Medical Prescriptions

In this Section, the use of the proposed platform featuring the CONVERGENCE architecture is described in the context of a scenario regarding the prescription of medical supplies, such as medicines. The scenario described tries to demonstrate the added value that CONVERGENCE can offer as compared to the existing solutions as regards searching, retrieving and recommending information in a safe, secure and accurate way. Identity, location and time aspects of the context-awareness are enhanced in order to overcome the various sensitive issues that such an application has to deal with. In the following paragraphs, we will elaborate on the scenario, starting from the identification of a real life problem in which the deployment of the platform can provide a solution: that of medical prescriptions. It should be noted here that the use case is conceptual, and has as a goal to focus on the potential of the platform for offering a universal way for fast and secure access to information that is linked to physical objects.

It is well known that a medical product can have different naming, depending on the country of export or production. Therefore, when travelling abroad, people who are taking medicines, or just as a measure of precaution, usually carry with them sufficient supplies for the whole period. This behaviour may avoid any difficulties they could have in purchasing locally the required medicines, due to the probable un-ability of the pharmacist in determining the corresponding local name or an equivalent medicine. But this is even the beginning, as there are far more serious issues involved with medicine prescriptions: side effects are quite common among patients that receive prescriptions, while incompatibilities between substances and the special conditions of patients may even lead to dangerous situations. In the proposed platform structure, each medicine has a unique VDI linked to it, where all relevant information is stored. This information is provided by the manufacturer of the product and is unique for every version of the product throughout the world.

Following these requirements, each pharmaceutical company is responsible for creating and updating a VDI for a specific product including information such as ingredients and compatibility list, side effects, precautions and prohibitions during the treatment. The company is responsible for updating the VDI, in cases of price or availability fluctuations, as well as withdrawal, or substitution. The National Health Organization (NHO) of a country or the Ministry of Health (MH) are responsible for validating each separate medicine VDI and ensure that all medicines contained in the platform are not harmful in any way to the general public.

Based on this information, doctors can make use of the ontologies created for both the medicines and the ingredients, in order to select alternative treatments for patients or exclude medicines that may be dangerous to some patients (i.e. pregnant women). The fact that VDIs related to health conditions (pregnancy, lack of enzymes), medical history and special conditions for each patient exist, can be used by the doctors in order to support the medicine selection, and even exclude medicines that are inappropriate for specific patients. This process can even be performed without complete access to the medical record of the patient, as comparison and compatibility of the related fields of the medicine and patient VDIs can be done with respect to the protection of privacy, while access to the patient's VDI can be granted to the doctor for specific purposes and under restricted conditions.

This can be easily supported through the inclusion of personal devices with geolocation capabilities, (i.e. mobile phones equipped with GPS), so that the user, through his personal device, can provide access to the doctor to specific parts of his personal health record. Exploiting context awareness, access is provided only when the patient's location is near the registered location of a certified doctor's office (available in a Central Database of certified doctors) and for a specific time period that corresponds to the visit duration (i.e. 1 hour). Upon provision of the prescription to the patient, the latter will be able to make use of the platform in order to safely get access to the medication (eliminating the case of asking for a different product from the pharmacy, which is important for small children and elderly people). Moreover, in case the pharmacy does not provide the prescribed medicine, for example because the patient is abroad and the medicine has a different name, the context-aware application can identify the patient's location and, using the underlying CONVERGENCE system, can suggest alternatives.

Furthermore, in the case where orders regarding general public health are issued (i.e. withdrawal of a product) this can be safely diffused to doctors and pharmacists, allowing for immediate actions towards protecting the general public.

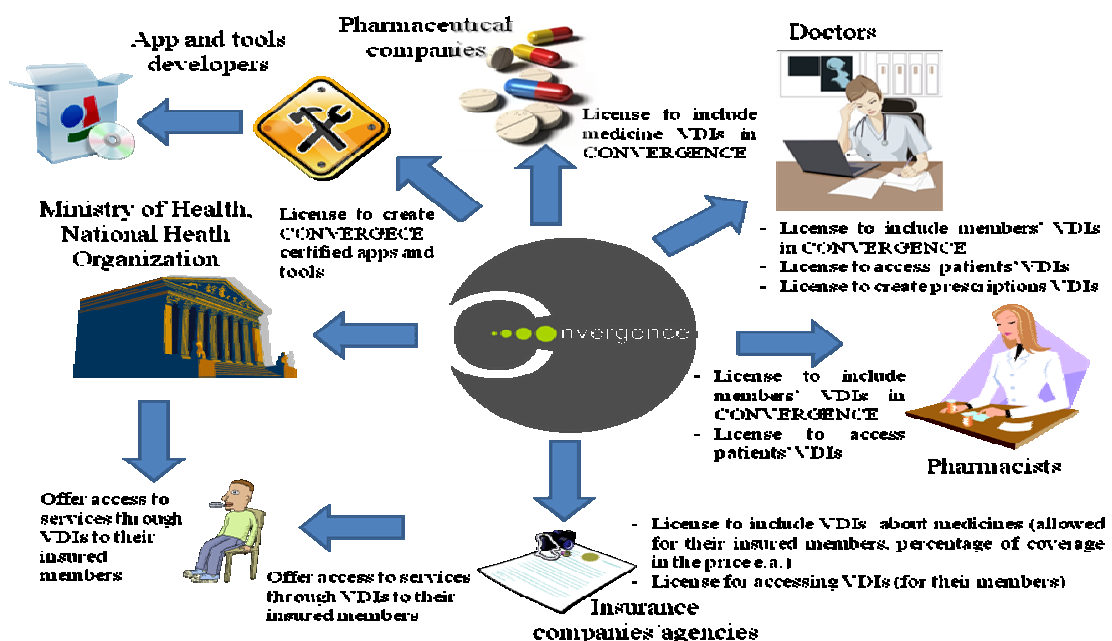


Figure 2: Medical VDI Platform Basic Interactions

4. A Deployment Scenario

The schematic diagram of medical information retrieval with the use of the proposed platform is shown in Figure 2. The CONVERGENCE system provider sells licenses to all involved members of the medical community. Pharmaceutical companies may wish to

publish VDIs concerning their products. These VDIs can be accessible from patients who wish to retrieve information on specific medicines as well as from pharmacists and insurance companies who wish to facilitate their clients. When a medicine is requested for a mobile user (i.e. a patient), either for writing or for executing a prescription, the user establishes connection with CONVERGENCE through his mobile device. At this point, it is assumed that this user is already registered to the medical VDI database, and therefore a VDI corresponding to his Personal Electronic Health Record (PEHR) is available. Provision for characterization of the information contained in the VDI in order to provide protection of personal information is important here. Using, from the one side, the VDI medicines database, which NHO or MH has verified, and, on the other side, the CDS, the system retrieves information on the prescribed medicines, concerning their suitability for the patient. Privacy is ensured, as no records of the VDI search are kept in the network after the search procedure. The CONVERGENCE application layer will then be able to perform a combined search in the medical record specific VDI and the medicines VDIs in order to help the doctor define the indicative treatment for the patient.

Finally, applications will provide the high-level associated functionalities related with medical VDIs, such as the ability to update or modify a medical VDI in cases of price or availability fluctuation, as well as medical VDI revoke if a medicine is verified to be harmful for patients. In turn, this application-level functionality is supported by the CoMid via its generic API, which enables applications to handle any kind of VDIs. Moreover, interoperability with context-aware applications will be considered as well, such as connection with GPS devices in order to route users to the medicine's sales location.

5. Conclusions

A content-centric and context-aware future Internet platform has been presented in this paper that facilitates medical information retrieval. This platform can be used either by patients who wish to retrieve information on specific medicines, or from doctors who want to define alternate treatments for their patients. In the initial deployment of the platform, tests will be carried out with well-known medicines while in future stages interoperability with well-known context-aware applications such as GPS navigation will be supported.

Acknowledgements

This research was partially supported by the European Commission under contract FP7-ICT-2009-5 CONVERGENCE [1].

References

- [1] Official web site of the Convergence project, www.ict-convergence.eu
- [2] ISO/IEC TR 21000-1: MPEG-21 Multimedia Framework Part 2: Digital Item Declaration
- [3] Network layer solutions for a content-centric Internet. A. Detti, N. Blefari-Melazzi in "Trustworthy Internet", Springer 2010.
- [4] A data-oriented (and beyond) network architecture. T. Koponen, M. Chawla, B.G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica in procs. of ACM SIGCOMM'07, 2007.
- [5] Networking named content. V. Jacobson, D. K. Smetters, et al., Fifth ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT), 2009.
- [6] On the Cache-and-Forward Network Architecture L. Dong, H. Liu, Y. Zhang, S. Paul, D. Raychudhuri in procs. of IEEE International Conference on Communications 2009, ICC 2009.
- [7] An Introduction to Petname Systems, Mark Steigler, available at <http://www.skyhunter.com/marcs/petnames/IntroPetNames.html>.
- [8] The Many Faces of Publish-Subscribe, P.T. Eugster, P.A. Felber, R. Guerraoui, A. Kermarrec, in ACM Computing Surveys (CSUR), Volume 35, Issue 2, 2003.
- [9] On the Personalization of Personal Networks - Service Provision Based on User Profiles. I. G. Nikolakopoulos, C. Z. Patrikakis, A. Cimmino, M. Bauer and H. Olesen special issue of the Journal of Universal Computer Science (JUCS), entitled "Intelligent Environments and Services", vol 15, no 12, [pp 2353-2372], 2009.

- [10] Measuring privacy through entropy in context-aware mobile services. C. Patrikakis, A. Voulodimos, IEEE PERSASIVE COMPUTING vol. 6, issue 4 pp. 73-74 Oct-Dec 2007.