

Cognitive and context-aware assistive environments using future internet technologies

Charalampos Doukas · Nikos Fotiou ·
George C. Polyzos · Ilias Maglogiannis

Published online: 18 June 2013
© Springer-Verlag Berlin Heidelberg 2013

Abstract Future Internet (FI) technologies are introducing new ways of networking and cognitive data delivery. In this paper, the potential of FI-based architectures for enabling the context-aware content adaptation and specialized delivery of health related information in assistive environments is investigated. The proposed system utilizes the publish/subscribe internetworking (PSI) architecture, an information-oriented architecture built for the FI using the so-called publish/subscribe paradigm. Information is brought at the center of the approach, providing several advantages: flexibility, seamless information morphing and exploitation of context, access control, and security in general. In addition to an overview of the approach and its characteristics, this work also presents the implementation of a subset of an assistive environment, using Blackadder, PSI's prototype, and illustrates its potential with an emergency service scenario for the assistive healthcare domain.

Keywords Content adaptation · Context awareness · Assistive information delivery · Publish/subscribe networking

C. Doukas
University of the Aegean, Samos, Greece
e-mail: doukas@aegean.gr

N. Fotiou · G. C. Polyzos
Athens University of Economics & Business,
Athens, Greece
e-mail: fotiou@aueb.gr

G. C. Polyzos
e-mail: Polyzos@aueb.gr

I. Maglogiannis (✉)
University of Piraeus, Piraeus, Greece
e-mail: imaglo@unipi.gr

1 Introduction

The introduction of the pervasive healthcare paradigm has improved awareness toward the elderly, healthcare for inhabitants of remote, isolated, and underserved locations, and the need for constant medical supervision of chronic patients. In this context, advanced electronic healthcare services that were once available only at healthcare centers are now expected to be available through a communication network anytime, anyplace, and to anyone. A medical assistive environment takes advantage of pervasive and ubiquitous technologies for delivering the above services. The term “assistive” is used more generally, including not only assisting persons with known health problems, but also empowering any human to improve the quality of life by exploiting sensing and computer-based support at all levels. Thus, the type of information that is exchanged in assistive environments is characterized by diversity and the need for content adaptation.

In such an environment, the promise of the future Internet (FI) and the new concepts and architectures that it brings (such as distributed computing and ad hoc networking, cloud computing and elasticity in resources, the Internet of things) can fit perfectly, enabling the development of advanced assistive applications. In this work, the application of FI-related technologies and architectures in assistive environments, focusing on advanced networking and context-awareness issues, is discussed. The proposed methodology is based on a specific clean-slate information-oriented architecture, the publish/subscribe internetworking (PSI) architecture. PSI is a FI architecture which utilizes information as its building block. It is based on the so-called publish/subscribe paradigm, and it is designed to support effective and flexible information governance and dissemination. Information-centric architectures appear to

be ideal for pervasive healthcare environments as they achieve effective information collection, provision, processing, and governance [1]. Thus, in the remainder of this paper, PSI is considered as the reference architecture for the introduction of FI technologies in assistive environments. The PSI architecture was designed within the EU FP7 PSIRP¹ project and is being further developed within the EU FP7 PURSUIT² project.

The remainder of this paper is structured as follows: Sect. 2 gives an overview of adaptation and context-awareness requirements in assistive environments and briefly presents various FI architectures. Section 3 details the elements of the PSI architecture and Sect. 4 describes proposed extensions that implement the foreseen cognitive assistive environment. Section 5 presents an implementation of a subset of the proposed architecture along with an initial evaluation, and finally, Sect. 6 presents future work and concludes the paper.

2 Background and related work

2.1 Information delivery and context awareness in assistive environments

The success of an assistive environment relies on the proper, accurate and on-time delivery of information and data exchange among the involved parties (users, operators, sensors, etc.). The strict requirements, in addition to the complexity of assistive environments (e.g., various devices, data types, user requirements, different medical personnel involved), introduces several issues that need to be addressed:

1. Assistive environments have special requirements on content and information presentation (e.g., in the case of elderly and people with vision impairment).
2. Different presentation layers exist. For instance, patient data have to be coded in different formats for transmission and decoded, respectively, for a medical expert to assess it. Different device specifications also introduce issues in data interoperability for both collecting and transmitting various contextual data. Moreover, the access network of a user may be the bottleneck of a rich data transmission session.
3. Specialized delivery in the context of reactive/proactive data transmission (e.g., in cases of emergency) is needed.
4. The transmitted information usually contains sensitive data that need to be secured. Patient data should be

protected from eavesdropping, as well as from manipulation during transmission.

The following sections introduce the concept of context awareness in assistive environments and highlight the proposed system.

2.1.1 Context awareness

Context awareness is the capability of networking applications to be aware of the existence and characteristics of user activities and environments. In rapidly changing scenarios, such as the ones considered in the fields of mobile, pervasive, or ubiquitous computing, systems have to adapt their behavior based on the current conditions and the dynamicity of the environment they are immersed in [2]. A system is context aware if it can extract, interpret and use context information and adapt its functionality to the current context of use. The challenge for such systems lies in dealing with the complexity of capturing, representing and processing contextual data. In assistive environments, contextual information might refer to user context (e.g., special user requirements like nutrition, user status and location), the environment context (e.g., indoor or outdoor) and the hardware context (e.g., presentation and communication devices used, network status). To capture context information, some additional sensors and/or programs are required [3–5].

The way context-aware applications make use of context can be categorized into the following three classes: presenting information and services, executing an application, and tagging captured data. In more details:

- Presenting information and services refers to applications that present context information to the user, use context to propose appropriate selections of actions to the user or personalize information delivery.
- Automatically executing an application describes services or applications that trigger a command, or reconfigure the system on behalf of the user according to context changes.
- Attaching context information for later retrieval refers to applications that tag captured data with relevant context information.

One way to capture and distribute contextual data is through the use of intelligent agents. Intelligent agents can be viewed as autonomous software (or hardware) constructs that are proactively involved in achieving a predetermined task and at the same time react to their environment. According to [6], agents are capable of performing tasks (on behalf of users or other agents), interacting with users to receive instructions and give responses, operating autonomously without direct intervention by

¹ <http://www.psirp.org>.

² <http://www.fp7-pursuit.eu/>.

users, including monitoring the environment and acting upon it to bring about changes, and finally showing intelligence in order to interpret monitored events and make appropriate decisions. Agents can be as follows: proactive, in terms of being able to exhibit goal-directed behavior; reactive, being able to respond to changes of the environment, including detecting and communicating with other agents; or autonomous, making decisions and controlling their actions independently of others. Intelligent agents can be also considered as social entities, which can communicate with other agents, by using an agent communication language in the process of carrying out their tasks.

In the context of pervasive healthcare, intelligent agents can contribute by analyzing patient and contextual information, by distributing tasks to responsible individuals, and by informing users regarding special actions and circumstances. The necessary decision making is usually performed through data classification on the acquired patient signals and contextual information. Generated classification results can contain information concerning the status of a patient, suggested diagnosis, behavioral patterns, etc. Data classification is an important problem in a variety of engineering and scientific disciplines such as biology, psychology, medicine, marketing, computer vision and artificial intelligence [7]. Its main objective is to classify objects into a number of categories or classes. Depending on the application, these objects can be images, signal waveforms, or any type of measurements that need to be classified. Given a specific data feature, its classification may consist of one of the following two tasks: (a) supervised classification in which the input pattern is identified as a member of a predefined class; (b) unsupervised classification in which the pattern is assigned to a hitherto unknown class.

There is a vast array of established classification techniques, utilized in assisted environments (like those in [8]), ranging from classical statistical methods, such as linear and logistic regression, to neural network and tree-based techniques (e.g., feed-forward networks, which include multilayer perceptron, radial basis function networks, self-organizing map, or Kohonen networks [9], to the more recent support vector machines). Other types of hybrid intelligent systems are neuro-fuzzy adaptive systems, which can be comprised of an adaptive fuzzy controller and a network-based predictor. More information regarding data classification techniques can be found in [7].

2.1.2 Future internet architectures

The current Internet architecture has evolved around the so-called end-to-end principle. The Internet was originally designed based on the need of efficiently interconnecting mainframes and minicomputers and providing efficient

remote access to them. However, as the Internet evolved, users started focusing on accessing information items rather than remote machines; nevertheless, the core design of the Internet has remained unchanged. The host-centric model and the design based on the end-to-end principle of the current Internet architecture have been identified as the root cause of many of its limitations, including lack of effective mobility, multicast and quality of service support; security and economic issues add to the problems [10]. Various add-ons, such as network address translation (NATs), Mobile IP, content delivery networks (CDNs), peer-to-peer (P2P) overlays, all violate, in various ways, several aspects of the original Internet architecture in order to provide features that were not part of the original design (or the original requirements). Keeping all these in mind, various research efforts were launched in order to redesign the Internet architecture.

In this work, PSI is considered as the reference architecture and it is detailed in the next section. However, alternative architectures for the FI have been proposed in corresponding literature. Content-centric networking/name data networking (CCN/NDN) [11, 12] is an information-centric architecture that proposes routing based on—hierarchical—content/data names. In the specific architecture, consumers ask for content by broadcasting “Interest” packets that contain the name of the content requested. Any “Data” packet whose content name has the “Interest” name as a prefix is said to satisfy this interest (request). Note that the interest/data match in this architecture can take place in any network node, making difficult the definition of access control policies that are the core of our proposed system. 4WARD/SAIL [13] is another FP7 EU funded research effort that also advocates an information-centric Internet which will enable network diversity, allowing various types of networks to coexist and cooperate in a smooth and cost-efficient manner. The proposed architecture envisions the Internet as a network where subnetworks will be self-manageable and network paths will be active components that will be able to affect transport services. 4WARD/SAIL also intends to shed light on the socio-economic aspects of such an approach; therefore, its contributions could be supportive of our work. The data-oriented network architecture (DONA) [14, 15] and routing on flat labels (ROFL) [16] designs also explore the potential of using information as the core inter-networking component. DONA proposes a new identification scheme based on flat, self-certifying identifiers as a replacement for the DNS naming resolution scheme that enables “finding” and “fetching” content. ROFL investigates the possibility of having an internetworking architecture solely based on flat identifiers, using DHTs and hierarchical DHTs. The evaluation results of ROFL indicate that this approach is feasible and can incorporate all

the internetworking structures that exist in the current Internet. The PSI architecture borrows the naming concepts of these two projects, but it proposes a different way for organizing information, using scopes. Scopes are an essential component for the design and implementation of the proposed system and will be discussed below.

3 The PSI reference architecture

As already mentioned, in the context of this work, PSI is utilized as FI reference architecture. The PSI architecture is a clean-slate one, taking nothing for granted (not even the standard IP protocol and IP host-based addressing), rather applying a completely different communication paradigm at all layers of the architecture: the publish/subscribe (pub/sub) paradigm. Architectures that abide by this paradigm are mainly composed of three entities: the publishers, the subscribers and an event notification service [17]. Publishers advertise and provide information (content they have available). Subscribers are information consumers, who express their interest for specific information (content) by issuing subscription messages. The event notification service is responsible for matching the advertised content with subscribers' interests and for initiating a content forwarding process from publishers toward subscribers. The pub/sub paradigm allows for endpoints decoupling, as publication and subscription operations do not need to be synchronized (not even that publication precedes subscriptions). This decoupling offers significant security advantages and facilitates the deployment of mobility and multicast mechanisms [18]. Moreover, the pub/sub paradigm allows the existence of multiple publishers for a specific publication, enabling this way caching and multihoming [19].

Information is the building block of the PSI architecture; information is everything and everything is information [20]. Each information item in PSI is identified by a unique flat identifier named as the rendezvous identifier (RId). Information items are organized in scopes, which are physical or logical structures hierarchically organized that are used to locate information items, as well as to control their dissemination [21]. A corporate network is an example of a physical scope, whereas a (subset of members in a) social network is an example of a logical scope. Every scope is governed by dissemination policies, which dictate who can advertise information in this scope as well as who can subscribe to a scope's information items. Furthermore, every scope is identified by a unique flat identifier, named the scope identifier (SIId), and is managed by—at least—one network entity called the rendezvous node (RV). All RVs are organized in a network known as the rendezvous network. The rendezvous network enables scope lookup,

and it can be hierarchically organized (akin to DNS) or be flat (similar to DHTs). An information item may belong to multiple scopes. In order for a publisher to advertise an information item he possesses, he has to specify the item's RId as well as the SIId of the scope that will be responsible for this information item. The RId of an information item is application specific; it can be, for example, the result of a hash function over the item's data. Similarly, in order for a subscriber to receive an information item, a subscription message has to be issued toward the RV of the scope in which the information item belongs, while this RV becomes the so-called rendezvous point for the information item. Every publisher knows at least one RV, which is his default RV. So, as it is depicted in Fig. 1a, the publisher SERVER01 sends the advertisement message to its default RV that routes it to the appropriate RV. The specific procedure is illustrated in Fig. 1a, where the publisher is shown to advertise an information item to scope DE012, which is managed by RV05; it sends the advertisement message to his default RV, that is, RV01, and it eventually reaches RV05. When the advertisement message reaches the RV05, the publication list of the scope is updated. The subscriber therefore needs to know the SIId of the scope and the RId of the information item. SIIds and RIds can be learned by an out-of-band mechanism, such as a search engine (or direct communication). Upon a successful matching, the RV initiates a (multicast) delivery path creation process from a publisher toward the subscriber(s).

This process is achieved by utilizing information accumulated by a network of topology managers (TMs) and by taking into account subscriber requirements and possibly delivery policies. TMs are network nodes responsible for monitoring the network topology and for detecting topology changes. In order for an information item to be forwarded through a specific delivery path, a Bloom filter-based structure (a space-efficient probabilistic data structure that is used to test whether an element is a member of a set) can be used, known as the zFilter [22]. zFilters contain all the links that an information item needs to traverse in order to reach its destination(s). The zFilter formation process makes it difficult for a malicious entity to manipulate a delivery path. Figure 1b shows a subscription example. Subscriber CLIENT01 subscribes to the already advertised information item with RId A2 in scope DE012. When the subscription message reaches RV05, a forwarding path is created between SERVER01—which is the publisher—and CLIENT01. Finally, the publication is forwarded from SERVER01 to CLIENT01.

The PSI architecture seeks to secure the information itself rather than the communication channel. Therefore, secure information can be delivered even over unsecured channels. Scopes enable information access control, as policies can be defined over who can advertise information,




Fig. 1 **a** Publication (advertisement) and **b** subscription procedures in the PSI architecture

as well as over who can subscribe to advertised items. Information transmissions can be secured using packet-level authentication (PLA) mechanisms that guarantee the confidentiality and the integrity of the transmitted items. At the higher levels of the architecture, existing security solutions can be applied with minor modifications [23]. The pub/sub paradigm deployed by the PSI architecture offers a significant level of protection against (distributed) denial of service [(D)DoS] attacks and spam, as there is no information flow as long as the receiver has not expressed interest on a particular piece of information; that is, the receiver in a pub/sub architecture is able to instruct the network which pieces of information shall be delivered to it. Moreover, no information is requested from a publisher, unless the publisher has explicitly denoted its availability, that is, unless the publisher has advertised the availability of this particular piece of information. Finally, the pub/sub paradigm enables the development of anonymization mechanisms as publication and subscription operations are decoupled. It is also able to achieve better network availability through the deployment of caching, multicast and multihoming.

4 The proposed assistive environment

In this section, we describe an architecture that could form a foreseen assistive environment or system for cognitive health data delivery overlaid over a PSI network. The purpose of this architecture is twofold: by continually monitoring the user's health status, it delivers useful information, including doctor's prescription information and reminders, medical videos and medication alerts, while in case of an emergency, it provides immediate proactive information to medical personnel and alerts the responsible individuals (i.e., relatives, attending physicians). The proposed architecture is envisioned to provide new inter-networking components for the assistive system while existing

applications are expected to be compatible with this system with small modifications.

4.1 Architecture components

The proposed architecture is composed of the following modules: *Storage Devices*, *Patient Monitoring Equipment*, *Information Visualization Equipment* and *Content Transformation Nodes*. All these entities are interconnected over a PSI underlay network (Fig. 2).

The *Storage Devices* are network nodes that are used to store large amounts of data, including users' electronic health records, medical images, informative videos, physicians' contact information. These nodes have large storage capacity and achieve availability through replication.

The *Patient's Monitoring Equipment* is composed of devices that monitor a subject's health status and activity. Such devices are sensors deployed in or on the body, cameras, microphones medical equipment, etc. They are associated with the identity of a specific person using a variety of identification techniques, ranging from username to password login to smart cards or even biometrics.

The *Information Visualization Equipment* is the equipment used to view various information items such as vital signals, medical images, prescriptions and other data related to user activity and physical status. This equipment can be composed of various devices of different capabilities, such as HDTVs, PC monitors, tablets and mobile phones.

The content transformation nodes are network elements that are able to transform various content types into other types; for example, they are able to transcode video, resize an image or read aloud text (i.e., transform text to audio).

All these nodes act as publishers and subscribers in a PSI underlay network. The PSI network is composed of the *Rendezvous Network* and the network of *Topology Managers*. The rendezvous network is the network responsible for the architecture's core functionality, as it enables information matching, publisher location determination




Fig. 2 Main components of the proposed assistive architecture

and selection, creation of delivery paths and the application of access control mechanisms. Topology managers on the other hand are these network nodes that provide information to the rendezvous network regarding the current topology. Such information can be, for example, that a user is connected to the architecture through a mobile network, or that a link failed.

4.2 Functions

The basic functions of the architecture are as follows: *Scope Definition*, *Information Advertisement*, *Information Subscription* and *Information Forwarding*.

With the *Scope Definition* function, scopes are defined. Each scope has its own dissemination strategy, which is defined using an attribute-based access control scheme. Access to the information in each scope is based on user attributes and the context. Scopes can be used to advertise general-purpose information items: For example, a scope in which information regarding a specific disease is advertised, or special purpose information items, and for example, a scope in which information regarding a specific patient's condition is advertised. For the former, general-purpose access rules can be used, such as "all doctors can advertise, all users can subscribe," whereas for the latter fine-grained policies have to be defined, such as "only devices of patient X can advertise, only doctor Y can

subscribe." A scope's policy can be changed as needed, by taking into consideration the user's context. For example, the policy of a scope that provides the medical records of patient X can be "if the health status of patient X is normal, only attending doctor Y can subscribe; in case of emergency, the relatives of X as well as all doctors of hospital M can subscribe." A node in the rendezvous network is chosen in order to manage a scope. Although rendezvous nodes do not store the actual content, they shall implement all these (potentially complex) policies, including (potentially frequent) updates from user devices.

With the *Information Advertisement* function, the various components of the architecture advertise the information that they generate. This information is advertised to the appropriate scope, for instance to the scope that best reflects the desired dissemination strategy. Information advertisement contains an information identifier that will be used for referring to this information item, the information item's location (i.e., the location of the storage node in which this item has been stored), publisher credentials and context information, as well as metadata, such as information (MIME) type, author, size, expiration time, description. The advertisement message is sent to the node of the rendezvous network that manages the desired scope. In order to make an information item available to multiple scopes, multiple advertisements have to be sent to the respective rendezvous nodes. Upon receiving an information advertisement message, the rendezvous node uses the

publishers' credentials and context information in order to validate that the publisher is eligible to advertise information under this scope.

Within the *Information Subscription* procedure, access to an information item is requested. A subscription can have many forms; it can be, for example, a subscription to a specific information item, such as “an X-ray image,” or a subscription for a collection of items, such as “medical advice from my doctor.” Information subscription and information advertisement functions are decoupled; therefore, subscriptions for content that has not yet been generated can be sent. For example, a patient may subscribe to messages that notify him that it is time to take his medicine, or a doctor may subscribe to messages that notify her that one of her patients is in an emergency situation (and when her shift ends, the subscriptions can be transferred to whoever replaces her). It should be noted here that all these notification messages are information items, also generated and advertised by the appropriate users or devices to the appropriate scope with the appropriate RId. As an example, a scope “patient X” can be defined with SId “G235” and it can be agreed that all requests for help are advertised in this scope with RId “H2EF21”; all persons that are interested in receiving “request for help” messages should send a subscription for H2EF21 to the rendezvous node responsible for managing G235 (and the authorized subscriptions will be honored). The information subscription message contains the identifier (SId, RId) of the information item to which the user wishes to subscribe (wildcards can also be used), user credentials and context information, as well as various metadata, such as desired content format, desired content type. User credentials and context information are used by the rendezvous node in order to validate that the subscriber is legitimate and allowed to access the information item in question.

Information is forwarded upon a successful subscription/advertisement match, which triggers the creation of a forwarding path from information storage location toward the subscriber(s). The path creation is implemented using information provided by the topology managers, and it has as input the subscriber's requirements and context. Subscriber requirements may include content adaptation for specific networks or devices. In this case, the forwarding path will include content transformation nodes. Information forwarding can be reliable or unreliable, depending on the content type transferred as well as the subscriber's needs.

5 A use case demonstration: smart delivery of healthcare data in emergency cases

In this section, we describe the implementation of a system based on the proposed architecture in order to demonstrate

its feasibility. The implementation is tested over the PSI test bed, a network of nodes across Europe, utilizing a PSI prototype node, code-named Blackadder [24, 14]. The use case around the implementation is a system that monitors the vital signs and the physical context of the user and disseminates appropriate information and alerts in case of emergency, guarantying the appropriate levels of security and privacy through access control, as well as on-time delivery of critical information. In this usage scenario, patients are using wearable heart rate monitoring devices. The latter are generating ECG and heart rate signals on a regular basis, which are made available (i.e., they are advertised to a certain scope) through the PSI network. These signals are thus available to a list of (predefined) doctors, who are the patient's attending doctors (who subscribe to this item and are in scope, but the architecture makes it very easy for this list to be dynamic, based on context, with potentially elaborate conditions). Moreover, patient monitoring equipment is capable of recognizing certain cases of emergency, such as low heartbeat rate, or a heart failure. In this case, an alarm is generated and a nearby hospital (potentially according to various preference conditions) is notified. In addition to the recent signals, data and alarms, the system provides access to the medical history of the patient, in the form of an electronic health record, that can reside (potentially in parts) either with the user (e.g., stored in a carried personal mobile device) or at a care unit the patient has previously visited.

5.1 Components implementation

As already mentioned, our architecture is composed of storage devices, patient monitoring equipment, visualization equipment and content transformation nodes, all interconnected over a PSI underlay network. For our prototype implementation and tests, as basic storage devices and visualization equipment, we utilize regular PCs connected to a PSI test bed. The Storage Devices contain information about the patient's context, such as the medical record, and act as local repositories for storing the acquired biosignal and healthcare data. In addition to the basic storage devices, the system utilizes the functionality of remote nodes that expose information such as medication instructions through Web services. The patient monitoring equipment sends directly the collected data to the storage devices, which in turn make this data available to the visualization equipment through the PSI network. The content transformation nodes have been omitted from the implementation of the proposed system. In the following subsections, the patient monitoring equipment, the underlay network and the API, it provides are presented in more detail and it is shown how they are utilized for implementing the desired functionality.

5.1.1 Patient monitoring equipment

The equipment utilized in the prototype implementation for acquiring and monitoring patient biosignals is the following:

- A mobile 3-lead ECG device from Gibson Technologies³
- A Polar heartbeat chest strap⁴ capable of continuously monitoring the patient's heartbeat.
- An Arduino microcontroller⁵ for collecting the ECG signal and heartbeat. Arduino is an open-source electronics prototyping platform based on flexible, easy-to-use hardware and software. It supports a variety of extensions (shields) that provide additional functionality (e.g., collecting motion data) and networking capabilities (ZigBee, Bluetooth, WiFi, 3G/UMTS, etc.).
- An Android-based phone that forwards the collected information to the infrastructure.

The ECG interface is intended for use by the patients themselves and is easy to install. It features a pluggable interface, which can be directly connected to the Arduino board and may be used to provide more detailed information about the patient status when an abnormality is detected in the ECG signal or the heartbeat rate.

An appropriate Android application that receives the signal information and forwards it to the storage devices by invoking the appropriate Web Service calls has also been developed. In case of emergency, the application notifies a dedicated machine—acting as proxy—which in turn sends the appropriate message to the network.

5.1.2 Underlay network and API

Blackadder, the latest PSI prototype node implementation, is used in the proposed system. Several nodes across Europe are inter-linked using OpenVPN (over the existing Internet), forming in this way a large overlay network composing the PSI test bed. All nodes are identified by a unique NodeID, which is propagated throughout the network using an OSPF-like protocol. A dedicated node acts as the topology manager (TM) of the test bed and is responsible for monitoring the network graph and generating network paths (through several nodes at the PSI level). A separate node acts as the rendezvous point (RV) of the test bed. All advertisements and subscriptions are collected at the RV. In case of an advertisement and subscription match, the RV requests from the TM to create a delivery path. It then notifies the (appropriate) publisher to

send data through this path. The path creation process currently supports multicast as well as multiple publishers for the same publication. In the latter case, the TM chooses to use the publisher which is closer (as determined by a network metric, such as delay, or hop count) to the subscriber.

Blackadder is a piece of software that provides the basic functionality for enabling advertisement (publication) of information, subscription to information and notifications for subscriptions. Blackadder has been developed for the Linux OS using the Click modular router⁶ and its source code is already publically available.⁷ Blackadder provides an API that can be used to build complex applications, so it can be regarded as the basic building block of applications running over PSI.

In this work, we have extended this API to support the following functions⁸

- `publish_scope(scope_id, payload)`. This function creates a scope identified by the `scope_id`. The public key of the patient is used as `scope_id`. The payload contains the following elements: the access control definition file, as well as a digital signature over for the concatenation of `scope_id` and access control definition file. This signature is generated using the patient's private key.
- `publish_info(scope_id, info_id, payload)`. This function advertises an information item, which can be, for example, the identifier of an ECG file or a medical image. The item is advertised under the scope `scope_id`, which means all access rules defined in this scope are applied to this item. This item is identified by the `info_id`, which is the result of a hash function over its data. The payload may include metadata such as the type of the published data, information about its creation date, device used to generate it. `Scope_id`, `info_id` and `payload` are digitally signed using the patient's private key.
- `publish_data(scope_id, info_id, payload)`. This function publishes (forwards) the data connected with an already advertised `info_id`. The payload includes the following: the content to be published and a digital signature over the data, the `scope_id` and the `info_id`. A `publish_data()` API call is always invoked after the respective `publish_info()` call, that is, an information item should firstly be advertised—using `publish_info()`.
- `subscribe_scope(scope_id, payload)`. With this function, the subscribers subscribe to a scope. The payload contains the subscriber credentials, a digital signature

³ <http://www.gssted.com/R-ECG1.html>.

⁴ <http://www.polarusa.com>.

⁵ <http://www.arduino.cc>.

⁶ <http://www.read.cs.ucla.edu/click/click>.

⁷ <https://github.com/fp7-pursuit/blackadder>.

⁸ Each function call requires additional parameters—such as payload size—which are omitted for clarity.

over the credentials and the `scope_id`. With this function, the subscriber will learn all the items that are advertised in that particular scope. Moreover, this function creates persistent state, that is, the subscriber will be notified about all future items advertised in that scope. Nevertheless, in case of subscriber disconnection, these notifications will be lost. Separate mechanisms are provided in order to handle subscriber disconnection and mobility (e.g., a subscriber could subscribe from multiple locations). Subscriptions to scopes never expire; therefore, a leasing/renew mechanism is not necessary.

- `subscribe_info(scope_id, info_id, payload)`. This function enables subscription to a publication already advertised in the `scope_id`. The payload contains the subscriber's credentials, a digital signature over the credentials, the `scope_id` and the `info_id`. Since in the assistive architecture all information items are mutable, subscription messages do not create persistent state; every time a subscriber receives the desired information item, the state that the subscription message has created is erased.

For the `publish_scope()` and `publish_info()` functions, the relative `unpublish` function exists. Similarly, for each `subscribe_scope()` and `subscribe_info()` functions, the relative `unsubscribe` function exists.

An interesting feature of the proposed API—compared to the core Blackadder API—is that it does not introduce any constraint on the length and the form of the identifiers. This is particularly helpful as it allows the usage of (encryption) public keys as identifiers, which—as will be discussed in the next section—enables the secure communication of the various entities of the system.

Using this extended API, two higher layer RV nodes have been built: the first handles the scopes that are created by the patients and the second handles the scopes used for advertising emergency events (further elaboration on these two kinds of scopes is provided in the next section). The higher layer RVs use the Blackadder's topology management function in order to create paths among arbitrary nodes. It is noteworthy, however, that the prototype, in its current form, cannot handle dynamic topologies; the network graph has to be fed as a configuration file to Blackadder's TM.

5.1.3 Functionality of the system

In the described implementation, two types of scopes are distinguished: the first type is the patient's private scope, where patient's data are advertised, and the second type is a general-purpose scope, where emergency events are advertised. All patients have a public/private key pair. Each

patient's scope has his public key as ID—denoted as `Patient_PK`—whereas the emergency scope has a well-known pre-configured ID, which for the remainder of this paper is denoted as `EVENT_SID`.⁹ Moreover, two types of identities are used: the personal identity, which is assigned to a single user, for example Bob, and the role identity which is assigned (in general) to a group of users, for example doctors of ER. A public/private key pair identifies each single user and each specific role. The rendezvous point is provided as a service by the network. Users and devices do not have to be aware of the rendezvous implementation details, but they should be aware of its public key—denoted as `RVService_PK`—which is the key with which they should encrypt the payload of all the messages that are sent to the rendezvous system. In the current implementation, storage devices are owned by the patient; thus, the patient's private key is known to them.

Initially, each patient creates his private scope by calling the `publish_scope(Patient_PK, payload)` function. In the payload, for the current static implementation, the patient's agent inserts two lists of public keys. The first list is the list of persons that have access to his medical data, whereas the second list is the list of persons and/or roles that have access to his medical data in case of emergency. In case a patient wants to modify the access control lists, he simply has to call again the same function with the new access control lists. The payload is encrypted with the `RVService_PK`.

Doctors have to subscribe to their patients' scope (through their terminal devices). This is achieved by executing the `subscribe_scope(Patient_PK, payload)` function, once for each patient, providing their credentials in the payload. If a doctor wants to stop receiving data from a specific patient, the `unsubscribe_scope(Patient_PK, payload)` function can be used. The payload is encrypted with the `RVService_PK`.

Every time a storage device receives new data, it calculates its ID by applying a hash function. Then, it executes the `publish_info(Patient_PK, info_id, payload)` function. Almost immediately, the doctors in the access control list who have subscribed to the patient's scope will receive a notification that new data are available. If a doctor is interested in receiving the advertised data, she (i.e., her software agent on her terminal device) uses the `subscribe_info(Patient_PK, info_id, payload)` function. As a next step, the storage device receives a notification that there are available subscribers. Upon receiving such notification, the storage device sends the data by executing only once the `publish_data(Patient_PK, info_id, payload)` function and all the doctors that are in the access control

⁹ It should be noted here that other general-purpose scopes can be created for other events. Moreover, different types of emergency events can have their own scopes.




Fig. 3 Example of data exchange under normal conditions

list and have subscribed to this data will receive it. This is achieved because the TM has already calculated and configured the delivery tree from the storage devices to doctor's visualization equipment. Every time a storage device erases some data, it executes the `unpublish_data` (`Patient_PK`, `info_id`, `payload`). Figure 3 shows an example of data exchange under normal circumstances.¹⁰ In this example, a patient (`Patient_A`) wears a heartbeat chest strap. This strap generates ECG signals, which are sent to two out of the three storage devices using `HTTP_POST`. If a scope for the patient has not been created, one of the storage devices executes the `publish_scope` function. All the storage devices that receive the ECG data advertise it to the PSI underlay network. At the same time, a doctor, `Doctor1`, subscribes from his terminal device to the patient's scope. The PSI underlay network assures that `Doctor1` is allowed to subscribe to this scope and notifies the doctor's terminal about the new ECG signal. As a next step, the doctor subscribes to receive this specific data. The PSI underlay network through the TM calculates a path from the storage device, which is closer to the doctor's terminal, which in this case is the `Storage Device 3`. The underlay network notifies `Storage Device 3`, informing it of the path identifier. Finally, `Storage Device 3` publishes the ECG, which is forwarded to the doctor's terminal through the created path.

In case of emergency, the patient advertises an alert in the emergency scope using as `info_id` his public key; therefore, the following function is executed: `publish_info` (`EVENT_SID`, `Patient_PK`, `payload`). All the rendezvous nodes that are managing patient data will receive this notification. Then, the rendezvous nodes will relax the access control policies that control patient data and will notify the subscribers that are included in the emergency access control list about the data that are available in the patient's scope, providing at the same time information about the emergency event. The subscribers will examine

this data and will subscribe to the data item(s) that are necessary in order to handle this emergency event. When the patient issues an `unpublish_info` (`EVENT_SID`, `Patient_PK`, `payload`), the users included in the emergency access control list will not have access to this patient's data any more. Figure 4 gives an example of emergency case scenario. `Patient_A` carries a heartbeat chest strap which captures bio-signals and detects patient status. Initially, the rendezvous node that manages the patient's biosignals subscribes to a special scope used for emergency notifications. Moreover, the healthcare unit (i.e., `Hospital1`) that is responsible for `Patient_A` in case of emergency subscribes to receive his medical data. In this example, the monitoring equipment detects in the captured ECG signals a heart failure; therefore, it calls a Web service located in `Storage Device 1`. This device advertises an event in the emergency scope, which eventually notifies the rendezvous system that `Patient_A` is experiencing a heart attack. This notification results in the relaxation of the access control policies and the notification of the `Hospital Terminal1`. The notification message contains the `Patient_A` identifier, a code corresponding to the type of emergency, as well as a list of the IDs of the advertised data. In this case, the doctor monitoring this terminal decides to subscribe to a particular ECG file, `cardio_1`. This subscription leads to the creation of a forwarding path from `Storage Device 2`, which is the storage location for this specific data, to the terminal.

5.2 Initial evaluation of the system

In this section, we evaluate the proposed system, mainly from a qualitative perspective, in terms of performance, extensibility and security.

5.2.1 Scalability and performance

The developed prototype allows the deployment of multiple storage points and supports replication and multicast.

¹⁰ For clarity, payloads have been omitted.




Fig. 4 Messages exchanged in case of emergency

Each patient is not limited to a single storage point, and multiple devices can be used for the same patient. Being location independent, the proper operation of the developed architecture depends only on the IDs assigned to the information items. The generated ECG files and vital data can be replicated to multiple storage devices deployed in various parts of the network. When the TM is requested to generate paths between storage devices (which act as publishers) and doctor's visualization equipment (which act as subscribers), the storage device that is closer to the doctor's terminal, in terms of (some) network distance, will be selected. Moreover, if multiple doctors request the same item (and this is a typical case as all the doctors of a patient may have subscribed to receive everything that is advertised in the patient's scope), the TM will create a multicast delivery tree from the appropriate storage device toward the physician's visualization equipment. Every request in the system concerns a unique object identifier and not a network location. Requests on object identifiers facilitate the deployment of caches.

A system based on the current Internet architecture would require CDN-like add-on solutions in order to allow content replication and optimal selection of content sources, since the current Internet architecture does not natively provide caching, nor multihoming. Moreover, since the current Internet architecture does not support multicast in an effective way, the application developers should also be concerned with adding application layer multicast support in their systems. In the current Internet, all operations refer to network locations rather than objects. Therefore, in order to achieve network caching, application layer or deep packet inspection mechanisms have to be deployed.

5.2.2 Extensibility

New devices and users can be added to the developed system in an easy and straightforward manner. Any device equipped with the proper API can interact with the PSI network, with little or no configuration. It should be noted

that this API is not system specific; on the contrary, it is a general-purpose API provided by the underlay architecture, which will be used by other systems as well. This way, interoperability across applications is better supported. In the current Internet, system-specific mechanisms, such as Web services, have to be developed in order to support device and application interoperability. Whenever a new user is added to the system, there is no need for configuration in the rendezvous system, which is responsible for applying the access control policies. Patients wishing to enter the system have simply to create a new scope, which uses their public key as ID. The rendezvous system is able to determine that a new patient is a valid user by verifying their public key using the digital signature of the packet. In order for doctors to enter the system, they have to simply send a subscription message with their public key. The rendezvous point is able to determine the validity of the subscription by verifying the public key used in order to sign the message. If doctors are included in any of the patient's access control lists, they will receive the appropriate data. Forwarding in PSI is based on source routing; therefore, a rendezvous node can control the nodes through which data will be transferred. Source routing makes the deployment of special purpose nodes, such as content transcoders, easier. When such nodes are introduced into the system, a rendezvous node can be simply configured to direct traffic through these nodes if necessary. Since PSI is location independent, in contrast to the current Internet, a patient or a doctor is not bound to a particular network location. So, if a patient decides, for example, to provide their medical records from a different storage point, the overall system operation will not be affected at all, or if a doctor changes office and logs on to the system from another terminal, their session will be transferred to the new terminal. In the current Internet architecture, this cannot be easily achieved; either all endpoints will have to be reconfigured in order to communicate with the new endpoint or indirection mechanisms—such as DNS—will have to be used (and many times in ways not intended).

5.2.3 Security

The exchanged messages are digitally signed with the sender's key; therefore, their integrity can be verified. Moreover, by having all the messages signed, effective accountability mechanisms can be deployed, at least at the application layer. Data confidentiality is achieved by encrypting the payload using the RVService_PK. All the notifications sent from the rendezvous system to the publishers are encrypted using their public key. Each item is identified by the public key of the owner, along with the result of a hash function applied over its data. Upon receiving an information item, a subscriber is able to verify its provenance (by evaluating the digital signature of the message using the publisher's public key) and its relevance (by evaluating the hash of the received content). Moreover, since data are encrypted, it can be transmitted over an unsecured channel. Information item verification can also be done by in-network mechanisms. In case of multicasting, a different solution should be considered. In this case, the rendezvous system creates a group communication key, which is sent to the subscribers and to the publisher. This solution obviously requires a side channel in order to send the key.

Access to patient data is assured using access control. Provided that the rendezvous system is reliable, only legitimate users can access patient files. Moreover, removal (revocation) of a user from an access control list, or addition of a new user, is straightforward as a single message is required (containing the new access control list).

On the other hand, security in the current Internet is closely tied to data location (e.g., security certificates are bound to a specific URL) [25]. Moreover, security in the current Internet mainly concerns securing the underlying communication channel (e.g., using HTTP over SSL/TLS in order to provide secure HTTP). The proposed architecture relaxes these limitations. In order for an application that uses the current Internet architecture to achieve the security properties of the system, application-specific mechanisms need to be developed.

5.2.4 Cost

A concern that may arise regarding the implementation and installation of the proposed system is its cost. Nevertheless, providing that the underlay FI architecture is deployed, the cost for adapting existing equipment and applications is expected to be low. In our implementation, proxies have been used, translating HTTP messages sent from the patient equipment into PSI-specific API calls. Such proxies could be used in a real, incremental deployment. However, since proxies may introduce performance penalties, the

gradual modification of the communication protocol of critical equipment should be considered.

The previously analyzed evaluation is summarized in Table 1.

6 Discussion and conclusions

Context-aware assistive environments face difficulties due to limitations of the current Internet technologies. 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 producing and retrieving information and in particular because the generated content includes sensitive data, which has to be properly secured and kept private.

In this paper, a conceptual solution is presented for a context-aware assistive environment based on a future Internet architecture: the PSI architecture. In addition, it is shown that by designing an architecture around information items and by providing functions for manipulating and

Table 1 Comparison between the proposed system and a similar system that uses current internet technologies

	PSI-based assistive system	Current Internet-based assistive system
Scalability and performance	Support for multihoming Support for multicast Caching can be easily deployed	Application layer solutions required for effective data transmission (such as CDNs, application layer multicast) Deep packet inspection for caching, or application layer caching
Extensibility	Better support for interoperability Source routing facilitates the introduction of special purpose in-network nodes Location-independent mechanisms facilitate the mobility of the various stakeholders	System-specific mechanisms—such as web services—have to be created in order to support system interoperability Separate indirection mechanisms are required in order to support user mobility
Security	Information-oriented security mechanisms Information can be transmitted over unsecured channels Self-certifying names on information items	Location-based mechanisms application layer solutions have to be deployed in order to secure content (such as encryption key exchange mechanisms)

securing them, pervasive health assistive solutions can be deployed more easily and effectively. Moreover, part of this conceptual architecture has been implemented by utilizing a PSI prototype. Although in small scale, the implementation gives a hint about the scalability, extensibility and security features of such architectures and systems.

As the PSI prototype evolves, it is expected that functionality currently implemented by add-on solutions will be incorporated into the system's core design. In the current prototype, a single machine performs the rendezvous function. Future releases of the prototype will include distributed rendezvous functionality enabling more scalable handling of publications and subscriptions. Currently, the topology manager of the prototype creates paths among endpoints 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 policy requirements, for example, nodes through which data should be sent or nodes and networks to be avoided. This new feature will also enable the deployment and exploitation of content transformation nodes: Subscribers or publishers requiring content transformation will request the delivery paths to include these nodes.

In the adopted security approach, public keys were used to define the desired access control policies. Future work in this domain includes the possibility of the definition of access control policies using additional attributes. Such attributes can include advanced location and context-aware policies. For example, policies, such as “in case of emergency, all nearby hospitals must be notified,” should be possible to be easily implemented (and in very dynamic fashion, for example, considering not only the location of the emergency, but also time of day, traffic conditions, other emergencies). The API produced during 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, containment and access to information.

It should be noted that currently, patient monitoring devices are communicating directly with the storage devices owned by the patient using HTTP. However, the intention of the architecture and future work plans include allowing the devices to interact with the core PSI network. If user devices are directly integrated into the network, then storage can be regarded as a network service, as cloud storage is provided today, or in-network caches will be in the near future.

Finally, future work plans include the implementation of an overlay¹¹ version of the system, which will allow coexistence with IP networks and thus a faster path to

adoption through backward compatibility. This overlay version will also enable extensive, large-scale tests with real users and data.

References

1. Trossen, D., Pavel, D., Guild, K., Bacon, J., Singh, J.: Information-centric pervasive healthcare platforms. In: Proceedings 4th International Conference on Pervasive Computing Technologies for Healthcare (Pervasive Health) (2010)
2. Khedo, K.K.: Context-aware systems for mobile and ubiquitous networks. In: Proceedings International Conference on Networking, International Conference on Systems and International Conference on Mobile Communications and Learning Technologies (ICN/ICONS/MCL) (2006)
3. Doukas, C., Maglogiannis, I., Kormentzas, G.: Advanced telemedicine services through context-aware medical networks. In: Proceedings International IEEE EMBS Special Topic Conference on Information Technology Applications in Biomedicine (2006)
4. Maglogiannis, I.: Introducing intelligence in electronic healthcare systems: state of the art and future trends lecture notes in computer science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 5640 LNAI, pp. 71–90 (2009)
5. Doukas, C., Maglogiannis, I., Karpouzis, K.: Context-aware medical content adaptation through semantic representation and rules evaluation Proceedings—3rd International Workshop on Semantic Media Adaptation and Personalization, SMAP 2008, art. no. 4724861, pp. 128–134 (2008)
6. Fox, J., Beveridge, M., Glasspool, D.: Understanding intelligent agents: analysis and synthesis. *AI Commun.* **16**(3), 139–152 (2003). IOS Press
7. Zhai, J.H., Zhang, S.F., Wang, X.Z.: An overview of pattern classification methodologies. In: Proceedings of the Fifth International Conference on Machine Learning and Cybernetics 3222–3227 (2006)
8. Doukas, C., Maglogiannis, I.: An assistive environment for improving human safety utilizing advanced sound and motion data classification. *Univ. Access Inf. Soc.* **10**, 217–228 (2011)
9. Witten, I.H., Frank, E.: Data mining: practical machine learning tools and techniques, 2nd edn. Morgan Kaufmann, San Francisco (2005)
10. Feldmann, A.: Internet clean-slate design: what and why? *ACM SIGCOMM Comput. Commun. Rev.* **37**(3), 59–64 (2007)
11. Jacobson, V., Mosko, M., Smetters, D., Garcia-Luna-Aceves, J.J.: Content-centric networking. Whitepaper. Palo Alto Research Center, Palo Alto (2007)
12. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L.: Networking named content. In: Proceedings of ACM CoNEXT, (2009)
13. SAIL Project: D-3.1 (D-B.1) The network of information: architecture and application. (2011) http://www.sail-project.eu/wp-content/uploads/2011/08/SAIL_DB1_v1_0_final-Public.pdf
14. Trossen, D., Parisi, G.: Designing and realizing an information-centric internet. *IEEE Commun. Mag.* **50**(7):60–67, (2012), special issue on ‘Information-Centric Networking’
15. Koponen, T., Chawla, M., Chun, B., Ermolinskiy, A., Kim, K.H., Shenker, S., Stoica, I.: A data-oriented (and beyond) network architecture. *ACM SIGCOMM Comput. Commun. Rev.* **37**(4), 181–192 (2007)
16. Caesar, M., Condie, T., Kannan, J., Lakshminarayanan, K., Stoica, I.: ROFL: routing on flat labels. *ACM SIGCOMM Comput. Commun. Rev.* **36**(4), 363–374 (2006)

¹¹ As in [19] and also considered as a deployment strategy for PSI in PSIRP and PURSUIT [20].

17. Eugster, P.T., Felber, P., Guerraoui, R., Kermarrec, A.M.: The many faces of publish/subscribe. *ACM Comput. Surv.* **23**(2), 114–131 (2003)
18. Xylomenos, G., Vasilakos, X., Tsilopoulos, C., Siris, V.A., Polyzos, G.C.: Caching and mobility support in a publish-subscribe internet architecture. *IEEE Commun. Mag.* **50**(7):52–58, (2012), special issue on ‘Information-Centric Networking’
19. Katsaros, K., Xylomenos, G., Polyzos, G.C.: MultiCache: an overlay architecture for information-centric networking. *Comput. Netw.* **55**(4):936–947 (2011) Elsevier, special issue on ‘Architectures and Protocols for the Future Internet’
20. Tarkoma, S., (ed.) PSIRP deliverable 2.3, architecture definition, component descriptions, and requirements. (2010) <http://www.psirp.org/>
21. Fotiou, N., Trossen, D., Polyzos, G.C.: Illustrating a publish-subscribe internet architecture. *Telecommun. Syst.* **51**(4):233–245, (2012), Springer, Special Issue on ‘Future Internet Services and Architectures: Trends and Visions’
22. Jokela, P., Zahemszky, A., Rothenberg, C.E., Arianfar, S., Nikander, P.: LIPSIN: line speed publish/subscribe inter-networking. *ACM SIGCOMM Comput. Commun. Rev.* **39**(4), 195–206 (2009)
23. Nikander, P., Marias, G.F.: Understanding pure publish/subscribe cryptographic protocols. In: *Proceedings of the Sixteenth International Workshop on Security Protocols*. Cambridge, England (2008)
24. Kjällman, J., (ed.): PURSUIT deliverable 3.2, first lifecycle prototype implementation (d3.2). (2011) <http://www.fp7-pursuit.eu>
25. Lagutin, D., Tarkoma, S.: Cryptographic signatures on the network layer-an alternative to the ISP data retention. In: *Proceedings of IEEE Symposium on Computers and Communications*, Riccione, Italy (2010)