

Towards Consolidated Presence

(Invited Paper)

Manfred Hauswirth*, Jérôme Euzenat[†], Owen Friel[‡], Keith Griffin[‡], Pat Hession[‡], Brendan Jennings[§], Tudor Groza*, Siegfried Handschuh*, Ivana Podnar Zarko[¶], Axel Polleres* and Antoine Zimmermann*

*DERI, National University of Ireland, Galway, Ireland

[†]INRIA Grenoble Rhone-Alpes and Laboratoire d'Informatique de Grenoble, France

[‡]Cisco Systems, Galway, Ireland

[§]TSSG, Waterford Institute of Technology, Waterford, Ireland

[¶]Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia

Abstract—Presence management, i.e., the ability to automatically identify the status and availability of communication partners, is becoming an invaluable tool for collaboration in enterprise contexts. In this paper, we argue for efficient presence management by means of a holistic view of both physical context and virtual presence in online communication channels. We sketch the components for enabling presence as a service integrating both online information as well as physical sensors, discussing benefits, possible applications on top, and challenges of establishing such a service.

I. INTRODUCTION

Mobile workers and virtual teams are rapidly gaining importance in the knowledge economy. Being able to communicate efficiently with colleagues, customers, partners, suppliers and peers is essential in today's global enterprise workspaces. *Presence* is a key ingredient to delivering on this capability as it provides the ability to automatically identify the status and availability of communication partners and resources both in terms of physical context (e.g., location, ongoing meetings, booked resources), and virtual presence in online communication channels (e.g., IP telephony, video conferences or Internet messaging (IM)). The availability of *presence* optimises communication time and hence time to resolution, in turn driving productivity increase, customer satisfaction and business revenues.

A study conducted by Chadwick Martin Bailey (July 2008) [1] found that on a daily basis 40% of employees are unable to reach co-workers on the first try resulting in more than 20% of their employers experiencing a missed deadline or project delay on a weekly basis. As a result, businesses of all sizes are either deploying or are evaluating the deployment of unified communication infrastructures to facilitate such collaboration between employees and between business partners. *Presence* is one of the foundational components of unified communication experience and provides the ability to connect with colleagues on the first try by knowing their availability in advance.

The supposedly simple but central question of availability—or, more generally, *presence*—involves complicated technical questions. For example,

- virtual availability does not necessarily translate into actual availability or presence;

- correct determination of a person's or resource's availability requires the tight integration of various sources of virtual and physical presence, for example, calendar information with physical availability and monitoring of physical availability;
- there is not one single view on presence, but many, determined by the private policies of the monitored person or object, corporate policies and security and privacy policies, so that the same entity may have different presence for different requesters.

Having an integrated, consistent and correct view of *presence* has thus become a standard requirement in many application scenarios to guarantee productivity in modern business environments. A meeting scheduler using presence services could, for example, reschedule a meeting rather than having attendees to wait for late-comers, where a new meeting time could be proposed based on the geographical location and estimated arrival times along with virtual presence of participants. Another day-to-day example would be that if a person is busy in her/his office or offsite, a visitor may leave a request for meeting assigned to the geographical location which would be delivered to the addressee on arrival at the physical location rather than sending an immediate, possibly interrupting request. There is a virtually endless list of possible scenarios which all essentially depend on an integrated view of virtual and physical presence.

This paper aims at analysing current definitions and approaches for presence and proposes presence, or more concretely *Consolidated Presence*, as a first-class type of service to applications as part of an “Internet of Services” consuming information from the “Internet of Things” [2] as a core source of information integrated with presence information from corporate information systems, Web data, etc. Our goal is to provide an open and integrated view of presence as a service to users and application developers allowing them to easily integrate arbitrary sources of presence information through the use of open semantic standards which will be developed in the course of the project.

We aim at extending the narrow view of person-associated presence into a general concept of presence as the contextualised availability status of a person or resource. As presence and availability of persons and resources are very sensitive

issues and business-vital assets, the actual implementations need to follow a flexible approach to express arbitrary policies for enabling multi-faceted views of presence, and provide guidelines for ensuring privacy and protection of sensitive presence information in federation scenarios.

Our focus is explicitly on corporate environments with clear governing and enforceable policies rather than on open infrastructures such as the Web with the associated, complicated privacy issues. We target distributed settings within the same organisation, for example, a company with multiple units which are geographically dispersed in one or more buildings or places, and on federated access and exchange of presence information among a number of such organisations, governed by formally specified and verifiable access policies.

Consolidated Presence shall improve both intra-enterprise and inter-enterprise presence management systems. While the focus on the enterprise seems to imply closed world scenarios, we also need to take into account free mobility of users along with their personal policies and the use of open Internet-based presence sources which results in a mixed environment. In such settings sensitive context information needs to be encapsulated and tunnelled to the corporate presence management system before being shared, which simplifies the privacy and security concerns, provided that the users consent to the tracking of their presence information which is a reasonable assumption within corporate environments.

We motivate the need for an improved notion of presence by concrete scenarios (§II), which current standards cannot fully capture (§III). Thereafter, we introduce *Consolidated Presence* (§IV) and its requirements (§V), leading to a proposed roadmap for the development of such an advanced presence system (§VI).

II. SCENARIOS

This section provides a set of example scenarios in a fictitious enterprise environment which help ensuring the comprehensive coverage of the presence domain in enterprises.

FictInc is a successful SME operating world-wide both in terms of development and sales. The employees of the company are involved in a lot of travel activity and meetings frequently involve video conferencing. To minimise organisational overheads, the company has equipped its premises with state-of-the-art sensors and all employees have smart phones equipped with location/positioning hardware and software (RFID, GPS, WLAN positioning, Bluetooth, OpenBeacon or similar). Additionally, each employee has access to multiple online communication channels, including POTS, mobile phones, IP telephony, Skype, IM, etc. Access to the sensitive position information is managed by a rule-based system enables the company to define and enforce a global access policy (agreed by the employees in their work contract and in line with legislation). This can be further refined by users to express their privacy requirements.

Presence-enhanced meeting scheduler and communicator. Currently FictInc is working on releasing a new product which requires frequent meetings of the development teams

in Germany (led by Inge) and Ireland (led by Sean), and may often involve the sales and marketing department in Sweden (led by Mia) which has been picked as the test market for the new product. In this setting, scheduling meetings can be quite difficult and cumbersome, if not supported by a flexible presence management system.

For a final review of the product's beta release Inge organises a virtual meeting with Sean and the German lead developer Hans. Inge uses her presence-enhanced meeting scheduler to organise the meeting: she defines the list of meeting participants, sets meeting priority to high, preferred time, meeting duration and location, while the scheduler suggests a meeting slot according to current participants' calendars. Fifteen minutes before the meeting starts the presence system detects that Sean is still in a higher priority telephone conference and has indicated that the call will probably last for another half an hour. Hans receives an SMS on his smart phone as he is on the way to the office; Inge gets a short pop-up message on her screen while she is editing a presentation. Finally, Sean and Inge start the video conference but Hans is late. The presence system detects that he is on his way to the meeting room and informs Sean and Inge that Hans will arrive in five minutes. When Hans arrives to the meeting room, the meeting starts.

During the meeting, one of Sean's team members sends him an IM. According to Sean's policy, based on the profiles of the people Sean is meeting with and his preferences, the IM system shows him as busy for his team members and the message is blocked until the end of the meeting. In Germany, Inge's secretary Hilde gets a request by FictInc's CEO for a meeting. As Hilde sees that Inge is in a meeting, she uses the presence-enhanced communicator to send her a short note. The presence service determines that Inge's laptop is physically next to her in the meeting room and the system produces a non-intrusive alert on it. Inge responds that she is available right after the meeting. The CEO is provided with an estimated time for a call and a time slot is booked into his calendar.

In the meantime, Hilde needs Inge to sign some papers but can see that she is again in a meeting. The signatures are not urgent but should be done before Inge leaves in the evening. Thus Hilde instructs the presence-enhanced communicator to notify Inge of the signature request as soon as she leaves her office and Hilde leaves the building for a delivery. Upon leaving her office for a quick coffee, Inge gets a reminder and drops by Hilde's office to sign the papers.

When Hilde returns she finds the signed papers, but notices a flaw in the document which needs to be checked with Inge. However, the presence system shows her busy in her office despite no meeting scheduled. The reason for this is that a sales representative together with a new customer dropped by Inge's office to get first-hand information on the roll-out of the new product. The presence system can detect the two employees of FictInc in the room via their tags, and the customer via her visitors badge. Hilde can deduce that a meeting is going without further information as the configured policy does not allow her to access this privileged information. Therefore

she instructs the presence service to notify her when Inge is available.

Federated scenario. FictInc is a very successful company and expands globally. They have recently acquired another company ExampleComp in a different geographic location and need to consolidate their networks. FictInc and ExampleComp use presence service infrastructure from different technology vendors. Both presence services must work seamlessly within the single enterprise domain. ExampleComp users had a small, global sales team that were never in the office and relied exclusively on mobile carrier hosted/IMS communications and presence services. As the presence system of FictInc uses open, semantically described data standards and externalises the functionalities described above via service-oriented interfaces, this integration can be done quite fast.

FictInc also has a small set of strategic partners with whom it closely cooperates: DistrInc is the exclusive distributor for FictInc products in a specific business domain. To prepare for the launch of the new product in Sweden, Mia wants to schedule a meeting with Alain, her main DistrInc contact for product shipments and checks if he is available through FictInc's presence system. Unfortunately, Alain is not available, so Mia resorts to checking for the presence of any other available person with the same expertise in DistrInc and sees that Lisa is available. Mia calls Lisa and arranges the details about the shipment to Sweden.

Mia was able to do this because FictInc's and DistrInc's presence systems have been federated a while ago, providing a policy-based view on the presence systems of each company to the other company, which ensures the proper and secure externalisation of presence information. In fact, the federation was quite fast as the presence systems of both companies rely on open semantic presence standards and service-based access.

III. CURRENT DEFINITIONS OF PRESENCE AND THEIR LIMITATIONS

Currently, there exists a number of standards for modelling several aspects of the presence domain. In this section we briefly analyse them and discuss their shortcomings.

The IETF working group SIMPLE¹ published a set of standards (RFC) for presence and presence-related information systems. They define an abstract model of presence, a data model, several data formats and protocols, among those SIP² and XMPP³. In particular, the extensibility of the XMPP protocol enables representation and sharing of different context elements, such as current location and user activities. Unfortunately, most of the extensions are unsupported by the majority of IM tools, and thus, not interoperable.

Presence is defined in RFC 3856⁴ as: “*the ability, will-*

¹SIP for Instant Messaging and Presence Leveraging Extensions (simple), Internet Engineering Task Force (IETF) working group – <http://www.ietf.org/dyn/wg/charter/simple-charter.html>

²SIP: Session Initiation Protocol – <http://tools.ietf.org/html/rfc3261>

³Extensible Messaging and Presence Protocol – <http://xmpp.org>

⁴IETF - A Presence Event Package for the Session Initiation Protocol (SIP) – <http://www.ietf.org/rfc/rfc3856.txt>




Fig. 1. Different roles in a presence service

ingness, or desire to communicate across a set of devices”. The abstract presence model which serves as foundation for both the SIP presence and XMPP specifications is introduced in RFC 2778⁵. It defines presentities, objects that expose their presence state, and watchers, objects expressing standing interest in presence information related to a set of presentities. Two entities are introduced in the presence model to handle the flow of information between watchers and presentities: presence agent (PA) and presence user agent (PUA). PUA manipulates presence information for a presentity and multiple PUAs are possible per presentity/watcher. PA is a logical entity capable of accepting subscriptions, storing subscription state, and generating notifications when there are changes in presence. Therefore, a presentity is a “provider” of presence information, while a watcher is a “requester”. The flow of information between presentity and watcher is facilitated by a presence service, cf. Fig. 1.

The data model described in RFC 4479⁶ describes the additional components that have to be modelled in a presence service, such as the end users, the devices and the specific services. Concretely, the model is encoded in an XML format called Presence Information Data Format (PIDF)⁷, specified via an XML Schema. The format has been extended in various ways, including temporal information⁸, calendar or activity details⁹.

The above listed approaches for capturing presence information fail in one or several of the following aspects:

- 1) limited types of person-associated availability are considered, rather than providing an open solution which enables the contextualised integration of arbitrary sources of presence information, be they physical or virtual;
- 2) individual or corporate access policies are either not associated to presence and cannot be used to flexibly reveal presence information, or such policies do not have clear and open semantics which is required for automatic integration of presence data and “understanding” of presence information and policies;
- 3) solutions are typically custom-built and cumbersome to

⁵A Model for Presence and Instant Messaging – <http://www.ietf.org/rfc/rfc2778.txt>

⁶IETF – A Data Model for Presence – <http://www.ietf.org/rfc/rfc4479.txt>

⁷IETF – Presence Information Data Format (PIDF) – <http://www.ietf.org/rfc/rfc3863.txt>

⁸IETF – Timed Presence Extensions to the Presence Information Data Format (PIDF) to Indicate Status Information for Past and Future Time Intervals – <http://www.ietf.org/rfc/rfc4481.txt>

⁹IETF – RPID: Rich Presence Extensions to the Presence Information Data Format (PIDF) – <http://www.ietf.org/rfc/rfc4480.txt>

integrate into applications as presence is not externalised as a service to be used in service-oriented architectures as part of the Internet of Services;

- 4) standards that enable the exchange of presence information and policies and their enforcement across applications and between enterprises are lacking or cover a coarse-grained view of presence only.

In addition, several vendors offer presence solutions and existing protocols and standards cover low-level protocols for presence exchange and federation exhaustively. Nevertheless, they suffer from notable drawbacks:

- 1) *physical presence* is not taken into account;
- 2) only limited control over when/where/how presence is available (limited policy and context dependence);
- 3) profiles for requesters (*watchers*) and providers (*presentity*) of presence are very limited or missing;
- 4) the granularity of disclosing presence information is rather coarse and privacy is limited or lacking.

This can be reduced to two major deficiencies of existing platforms: on the one hand, support for the integration of new presence sources, especially when these provide *rich* presence or context information that has to be aggregated, and on the other hand their inflexibility in terms of serving and dispatching presence dependent on the context and in particular applicable policies.

Integration of new presence sources. Processing and aggregating data from various heterogeneous presence sources, be it raw sensors or data from enterprise information systems in a scalable fashion is only possible to a limited extent with established technologies, since scalability problems and heterogeneity problems have to be dealt with at the same time in a dynamic fashion. For instance, presence of a person may be determined by its calendar, location, but also by the co-location of other presentities (that might indicate that the person is currently in a meeting), determined by sensor readings that have to be processed in an efficient, scalable fashion.

Moreover, the inclusion of the presence of devices and resources other than persons as first-class presentities is usually not well-supported in current systems. For instance, current systems natively support to subscribe e.g. to the availability of a meeting room or other physical resource, or, respectively such inclusion was only possible in an ad hoc fashion.

Context/Policy dependent presence. While the trend goes towards *rich* presence (including location, mood or other information beyond a simple “presence state”), the policies governing disclosure of such presence nowadays still too often follow an all-or-nothing approach, that is either full presence is revealed or not, but there is no fine-grained policy control in place that allows to reveal the right presence to the right requester at the right time, such as,

-  Out of town (for customers)
-  In a meeting (for colleagues)
-  Available for urgent calls only (for Boss)

Further, even in systems which partially enable vendor-




Fig. 2. Consolidated Presence “Food Chain”

specific policy control, federation even in intra-enterprise scenarios is limited by the narrow interface that current presence standards provide, and policy control and negotiation in inter-enterprise scenarios remains a largely open problem. We aim to support the full food chain from raw data to consolidated presence and applications on top, as illustrated in Fig. 2.

Support for policy integration and handling of policies is poor in the existing protocols and standards. This is partially because standard extensions to hook in such policies are missing, but also due to limitations of the core functionalities of the standard protocols: for instance, SIP provides a framework for subscriptions of presence watchers, but after acceptance there is no built-in standardised mechanism to serve different presence states to different subscribers. Also, presence subscriptions are always limited to explicit presentities, whereas semantic subscriptions, depending on the presentity’s context or profile, such as the following are not foreseen in the core protocols, e.g:

“Let me know whenever a technician knowledgeable in Linux is available”

“Let me know Jane’s availability only if we are in the same building”

“Let me know the availability of a meeting room for four with a whiteboard”

These issues become more severe in federation scenarios, since in systems which partially enable vendor-specific policy control or richer semantic description of presentities and watchers, federation in intra-enterprise scenarios is limited by the narrow interfaces that current presence standards provide: policy control and negotiation in inter-enterprise scenarios, as well as semantic search and discovery remain key obstacles to interoperable next generation presence management systems.

Our goal is to ensure that new technologies and architectures developed for *Consolidated Presence* can ultimately benefit all end users, businesses and enterprises. This can be achieved by developing open, interoperable and standard-compliant tools, and in particular, by advancing the underlying enabling technologies which we deem crucial components in the development of next generation presence management

systems.

IV. TOWARDS CONSOLIDATED PRESENCE

We extend the model of *presence* (see Fig. 1) by enabling the physical world to play a role in the presence management system. By doing this, dynamic context can be determined, including physical location and activity, concretely captured by sensor networks. Disclosed information is controlled by personal policies as well as corporate level policies, such that only those watchers who conform to a certain profile can access specific presence information, at a given time, in a certain dynamic context, thus making profiles (or static context) a major part of the presence model.

The new concept of *consolidated presence*, proposed in this paper, enables a requester (the *watcher*) to be served a policy-governed, contextualised view on the availability of a provider (person or resource, i.e., the *presentity*) as shown in Fig. 3, integrating

- both the presentity's and watcher's physical presence
- both the presentity's and watcher's virtual presence, and
- policies of the presentity, governing (corporate, legal) policies, and other relevant policies.

This definition of presence complements to the large body of work in the area of *telepresence* which focuses on investigating presence from a cognitive point of view, rather than as a practical service for collaboration. Telepresence focuses on cognitive presence aspects such as making virtual/telepresence appear as real as possible. However, our focus is quite different, as we concentrate on monitoring, modelling and delivering presence from the viewpoint of enabling collaborative work applications which may or may not include telepresence.

a) *Presence in an enterprise context*: *Presence* is defined as “the willingness and ability of a user to communicate across a set of devices with other users on the network” [RFC 3856]. We extend this notion to include resources and devices themselves. A *presence service* is thus a system that accepts, stores, and distributes presence information to interested parties. Since the main goal of a presence service is to communicate presence information with respect to user availability and capability to communicate, it is often regarded as the “dial tone” of the 21st century. The notion of *rich presence* refers to an enhanced form of presence awareness in which participants can determine whether other users are online, for example in a unified communication system, and if so, observe to a limited extent what they are doing, their location, mood, and so on.

We enhance *rich presence* by *physical presence* (through sensor technology) and *semantic presence*. By *semantic presence* we mean the user's presence determined through advanced Semantic Web and Intranet Search technologies. The consolidation of presence information coming from fine grained sensors as well as user-supplied information, and its filtering through context and policy, requires a precise modelling of this information and well as search and inference capabilities. *Consolidated presence* will be provided in a unified, service-oriented manner to applications and end users.

b) *Presence in a Federated Architecture*: In terms of communicating consolidated presence within and between enterprises, users must be able to communicate with peers both within the user's enterprise and with other enterprises or with presentities/watchers hosted on public Internet services. This means that presence information needs to be integrated with existing deployed infrastructure within the enterprise (**intra-domain federation**) and between enterprises (**inter-domain federation: business-to-business federation**) and between enterprises and their consumer users (**inter-domain federation: business-to-consumer federation**). Enterprises generally enter into contracts with service providers to provide some or all of their communication services. For example, mobile carriers may provide bundled and/or hosted services to enterprises including Partial Domain Hosting, Presence, Instant Messaging and SMS-IM Gateways. For us this translates into a requirement for **inter-domain and intra-domain federation: mobile carrier federation**. Much work is currently being carried out by 3GPP to standardise some of these service offerings.

All of these types of federation have their specific requirements in terms of platforms to be integrated, privacy, and policies which need to be addressed. We will define a flexible architecture, formats and APIs which can be used for all types of federation scenarios. Fig. 4 shows a conceptual view of intra-domain federation. We will abstract away from the peculiarities of different, heterogeneous presence services used within enterprises in terms of their underlying information models, policy support facilities, storage and processing facilities for rich presence information. Our architecture shall cater for presence services to be able to interchange presence information, notwithstanding different underlying protocols (e.g., SIP, XMPP), models, or policies; presence information may be collected from any Data sources within the enterprise, be it virtual data sources, sensors, or mobile devices.




Fig. 4. Intra-domain federation

In **intra-domain federation**, standardisation bodies are just beginning to address the necessary architectural models [3]. The supposedly simple case of enabling users to communicate and share information is complicated by the fact that (i) users within the enterprise may be using different communication infrastructures from different vendors and that (ii) a single user within the enterprise may be using multiple




Fig. 3. Consolidated presence

different communication infrastructures simultaneously. Also users may have fixed devices, installed soft clients accessing on-premise applications, browser/thin clients accessing hosted applications, PDAs, mobile devices, etc. Access to presence information thus must be abstracted in a service-oriented way by suitable middleware. On the input side, presence information needs to be filtered as close to the edge of the network as possible to reduce load and support scalability.

For **inter-domain federation**, presence services must be open and extensible so that enterprise partners and specialist application providers can readily integrate with and enhance presence services via open APIs. Users must be able to communicate and share information with other users in different enterprises in the same way as in the intra-domain case but governed by different policies and users should have access to presence services from within the enterprise and from outside the enterprise via secure connections or via their mobile devices. Mobile devices and hosted solutions, e.g., Skype or from a mobile carrier, additionally require the presence system to support complicated deployment scenarios in a secure way—both intra-enterprise and inter-enterprise. Secure and policy controlled information and communication sharing must be pervasive throughout the enterprise and at all boundaries between the enterprise and external enterprises or consumer spaces. Fig. 5 shows a conceptual view of inter-domain federation with a mobile carrier.

Recently, enterprises are also looking to leverage the wealth of information that is available in their data centres to increase productivity and provide business value. More and more tools that are widespread in the consumer space are now gaining popularity within the enterprise and are rapidly becoming indispensable, productivity-enhancing business tools. Knowledge workers use a vast array of tools and applications throughout their working day including blogs, wikis, chat

rooms, video-blogs, micro-blogging, forums, teamspaces, etc. A user's activities in these on-line communities can provide a wealth of information about his expertise, availability and presence. Semantically annotating this information, integrating this information with traditional communications mechanisms and enabling sharing of this information both within and between enterprises has the potential to change the way that we work. By the adoption of Semantic Web standards for knowledge representation we will also enable the simple integration of this derived presence information into a unified model and infrastructure.

V. REQUIREMENTS

From a pragmatic perspective, there are a series of requirements to be met, in order to develop an open, unified consolidated presence model and infrastructure for federated enterprise environments:

Information integration covering the virtual, physical and social presence of people, objects and software entities. Formally specified, semantically rich information models, facilitating expressive and precise representation of concepts relating to availability and presence, are required. In addition, these models should also encapsulate personal and corporate policies and be easily instantiated as knowledge bases used by state-of-the-art semantic information processing techniques to form and share consolidated views of presence for people, objects and software entities.

Powerful and flexible semantic techniques gathering information via low-level stream processing, sensor middleware and publish/subscribe systems. It is essential to develop a middleware infrastructure that can flexibly access and integrate presence related information from a wide range of sources, including sensor networks, presence updates from software applications and activity traces scraped from online




Fig. 5. Inter-domain federation

sources. This will form a consolidated view of the presence of a person, object or software entity.

Enterprise policy management facilitating fine-grained control of the sharing of presence information by individuals, both within single enterprises and across enterprise boundaries. As the focus is mostly on enterprise environments, a proper enterprise-focused policy management solution needs to be developed. This should incorporate policy authoring tools that facilitate delegation of policy authoring capabilities to individuals, powerful policy analysis processes that ensure that authored policies are mutually consistent and in line with corporate goals, together with policy negotiation and alignment processes to manage the lifecycle of intra- and inter-enterprise federations.

VI. DEVELOPMENT DIRECTIONS

In order to support the development of the *Consolidated presence* food chain and the requirements listed in the previous section, one should look into advancing technologies from the following areas:

- “Raw” presence → **Data acquisition Middleware for Personal Devices, Publish/Subscribe middleware**
- “Digested” presence → **Semantic description of context models and policies, Policy analysis and negotiation**

In the following, let us review currently existing technologies in these areas and point out possible innovations.

Data acquisition Middleware for Personal Devices. Both fixed sensors as well as sensors on personal mobile devices are becoming an important tool for information management in networked enterprises. Personal devices enable local and remote access to personal information and they simplify the collaboration of co-workers by means of applications which query and manipulate this information, e.g., to schedule joint meetings or to assign tasks to co-workers. Additionally the present time, more and more personal devices are equipped with various kinds of physical sensors such as GPS receivers, accelerometers, microphones, etc. The resulting wealth of information accessible through these devices makes them a key hardware platform for the acquisition of presence-related contextual information such as the current user location or activity [4]. Besides, the use of personal mobile devices can also overcome some of the disadvantages of fixed sensing

infrastructures such as the associated maintenance costs for large-scale deployments. In the recent past, this has caused the development of numerous lightweight filtering and classification strategies for different types of sensors [5], [6], [7], [8]. Furthermore, it has spawned the development of several applications for different scenarios. Examples include the cooperative gathering of road conditions [9] as well as the localised classification of the user activity which can then be shared via an online social network [10].

The need of interconnecting sensors on the network level to enable integrated data processing requires a flexible middleware layer which abstracts from the underlying, heterogeneous sensor network technologies and supports fast and simple deployment and addition of new platforms, facilitates efficient distributed query processing and combination of sensor data, provides support for sensor mobility, and enables the dynamic adaptation of the system configuration during runtime with minimal effort. The Global Sensor Networks (GSN) middleware aims at addressing these goals [11]. Sgroi et al. [12] suggest basic abstractions, a standard set of services, and an API to free application developers from the details of the underlying sensor networks with the focus on systematic definition and classification of abstractions and services. Hourglass [13] provides an infrastructure for connecting sensor networks to applications and offers topic-based discovery and data-processing services. Like GSN it tries to hide internals of sensors from the user but focuses on maintaining quality of service of data streams. HiFi [14] provides hierarchical data stream query processing to acquire, filter, and aggregate data from multiple devices in a static environment. IrisNet [15] proposes a two-tier architecture consisting of sensing agents (SA) which collect and pre-process sensor data and organising agents (OA) which store sensor data in a hierarchical, distributed XML database modelled after the Internet DNS and supporting XPath queries.

As opposed to both these projects the data acquisition middleware needs to be far more generic in its ability to integrate a broad range of physical as well as virtual information sources providing various data types at a semantic level. A uniform model for integrating not only physical but also virtual information sources of presence information

is required, in order to enable their simple use in applications. The provided abstractions will facilitate the arbitrary combination of presence sources into higher-level aggregates which will be governed by sophisticated policies to determine use-case-driven presence assessment and presentation. The middleware should specifically take into account support for mobile devices to enable ad hoc collaboration among mobile devices in the vicinity to improve the amount and quality of presence information. Finally, it should provide an extensible and adaptive filtering and pre-processing framework to dynamically balance the trade-offs between the accuracy of the gathered information and the required energy for acquiring it.

Publish/subscribe middleware. Publish/subscribe middleware is a well-established communication infrastructure for dissemination of data from information sources, publishers, to information destinations, subscribers in distributed environments [16] and common presence protocols crucially rely on this paradigm. Publish/subscribe infrastructure represents a communication backbone for presence as it actively disseminates presence state updates to users. Users (*watchers*) are monitoring presence status of their colleagues and friends (*presentities*): Watcher interests (*subscriptions*) are therefore bound to presentities, while the infrastructure takes care of accepting and routing presence state updates (*notifications*) in real time. In particular, SIP Presence uses SUBSCRIBE and NOTIFY methods defined for SIP [17] to route presence subscriptions and state updates. XMPP defines a protocol extension for generic publish-subscribe functionality which enables XMPP entities to create topics and publish information at those topics to be broadcast further to all entities that have subscribed to the topics [18].

The publish/subscribe infrastructure for presence is currently topic-based and distributed. Even though presence status changes are assumed to be generated by events triggered only by human intervention occurring with a frequency on the order of seconds to hours, such status changes generate high signalling load both on the client-server interface (e.g. within radio access network) and within the presence server network. In particular, presence-related signalling in IMS-based networks introduces relevant non-scalable overhead especially when it comes to inter-domain SIP signalling [19]. The signalling problem becomes even more significant in context-aware environments where the frequency of presence-related publications will by far overwhelm the frequency of state changes in current presence systems. It is therefore vital to introduce content-based filtering capabilities and highly-efficient algorithms within the publish/subscribe infrastructure for presence to reduce inter/intradomain traffic and enable filtering of presence updates close to presentities.

The requirements on publish/subscribe middleware for presence on top of the Internet of Things are the following: first, topic-based solutions have to be replaced by content-based solutions that integrate fast and efficient matching algorithms (e.g. [20], [21]) to offer fine-grained filtering of presence information. As consolidated presence offers semantically rich presence models, these have to be mapped to less expressive

data models such that the processing time needed for matching of publications to subscriptions is minimised. Second, distributed solutions with efficient routing algorithms stemming from the publish/subscribe domain [22], [23] are needed as they are tailored to minimise the generated traffic associated to presence status exchange. Moreover, such algorithms have to support mobility across various networks, devices, and access points. Third, since presence environments for enterprises are governed by policies that have to be taken into account when designing filters for intra- and inter-domain routing of presence data, as well as to adjust the views on presence to end requestor, publish/subscribe middleware has to integrate a solution for policy-driven publish/subscribe matching and routing taking into account various corporate, security, and privacy policies. Therefore, the context-aware and policy-driven federated presence service represents a highly dynamic and challenging environment for the underlying publish/subscribe middleware which has to be carefully optimised to take into account specific requirements of real-world deployments.

Semantic description of context models and policies. Finer control of presence information requires characterisation of context, profiles and policies. This means, in particular, that the system is able to characterise the situation in which entities, either people or resources, are. There has been much work on context modelling in the fields of human computer communication, pervasive computing or artificial intelligence. They provide generally different context models, however the context description is usually tailored to the specific task to be carried out.

As other works have shown [24], [25], [26], ontologies are appropriate tools for defining context information because their use does not require exact match of information requirements with available information. Ontologies can thus be used as a way to define the elements that can be found in context representation and the context elements that are sought by applications.

Semantic Web technologies can be considered as a breed between knowledge representation and web technologies. They offer knowledge representation languages that are both expressive and open: two useful features for expressing contexts. In particular, openness allows the dynamic extension of ontologies and knowledge descriptions. Moreover, Semantic Web technologies come as standardised languages with available supporting tools. This lowers the barrier to their adoption.

To arrive at an interchangeable, extensible model of presence, context and policies, we suggest building on ontologies and Semantic Web technologies. The Web Ontology Language (OWL) [27] can be used to define standard vocabularies along with axioms governing presence, location, availability, profiles and policies. Existing ontologies are available on the Web, such as the Online Presence Ontology¹⁰, Geo¹¹, GeoNames¹²

¹⁰The Online Presence Ontology: <http://www.milanstankovic.org/opo/specs/>

¹¹WGS84 Geo Positioning: http://www.w3.org/2003/01/geo/wgs84_pos

¹²GeoNames: <http://www.geonames.org/>

both geolocation ontologies, the PIMO¹³ ontology from the Semantic Desktop (Nepomuk project¹⁴), which can link presence information to available information from documents and files on a local computer.

Policies. Policy management is widely seen as an appropriate paradigm to facilitate high-level, human-specified cognitive decision-making in system and security management. Policies are typically formulated as event-condition-actions rules. These rules are normally specified by system administrators using policy specification languages such as Ponder [28], Rei [29], KaoS [30] and XACML [31]. XACML (eXtensible Access Control Markup Language), which is standardised by OASIS¹⁵, is widely deployed in the industry to enforce security policies for communications between individuals and groups in enterprise deployments.

Although models of the semantics of policies are crucial for policy analysis the majority of policy languages were not designed with formal semantics, an exception of which being KaoS and REI, which are actually based on Description Logics. Some researchers have sought to bridge this gap by building semantic models of prominent languages. For example, Kolovski [32] built a comprehensive Description Logic (DL) based model of XACML policies that can be used for various analysis tasks. Barrett [33] specified a generic DL policy model that seeks to embody concepts common to most policy languages. Whilst these works have concentrated on models of policies in isolation, few have considered integrating policy models with models of the systems they govern, exceptions being Strassner et al. [34] who created an integrated UML model of policies and context, and REI for which an engine exists that, can partially reasons over REI policies and domain knowledge in RDF and OWL. Apart from such DL-based approaches, Protune¹⁶ (developed in the REVERSE EU project) uses a rule-based language and engine that allows reasoning about policies, mainly negotiation. Remarkably, rule-based languages seem to be closer in nature to policies than DL-based approaches and policy modelling is also a major use case of the recently standardised W3C Rule Interchange Format (RIF) [35] – a rules language interoperable with and extensible by OWL and RDF domain knowledge [36].

Policy analysis and negotiation. Although policy management has been a subject of research for more than fifteen years, the uptake of policy management systems by the industry has been slow. One of the main reasons for this is that sufficiently powerful and generic algorithms and processes to detect inconsistencies in the specification of policies have not been developed. Verlaenen et al. [37] provides a classification of policy analysis processes which includes: conflict analysis (does the behaviour specified by a policy contradict that of one or more deployed policies?), refinement (can a policy

be realised by a set of policies that relate more explicitly to the managed system?), dominance checking (if a policy is removed will the system behaviour change?) and optimisation (can a set of deployed policies be altered so that policy evaluation is more efficient?).

To date, most of the research literature has focused on the problem of policy conflict analysis. Lupu and Sloman [38] investigated conflicts of access control policies, which they defined as occurring when a set of simultaneously applicable policies result in multiple decisions being equally applicable.

More recent work has investigated the use of system models to aid in conflict analysis. Davy et al. [39] use the DEN-ng telecommunications UML-based information model to develop an application independent conflict analysis approach in which relationships that may signify conflict are identified based on information itself encoded in the model. Researchers are also increasingly investigating the use of Description Logic models as a knowledge base for conflict analysis; examples include the work of Uszok et al. [30] or Lin et al. [40].

All the existing published works on policy analysis target management of systems by a single organisation that has authority to define behaviour for all of the relevant managed devices and for all human interactions with those devices. However, *Consolidated presence* explicitly targets policy-based management of communications in federations of enterprises. This includes the development of policy analysis processes that ensure consistency between an enterprise's own policies and the policies they have agreed with other enterprises to govern interactions between individuals partaking in cross-enterprise project teams.

VII. CONCLUSION

We have presented a holistic definition of presence – which we call *Consolidated Presence* – characterised by combining presence context from both virtual and physical sources, coupled with the enforcement of personal and organisational policies that ensure privacy within and across enterprises. While not yet available in current presence management services, we have sketched scenarios that could benefit from such a consolidated presence model and have presented requirements as well as a technology roadmap towards building a working infrastructure supporting consolidated presence in federated enterprise environments. Basic building blocks include sensor technology, content-based publish/subscribe middleware, semantic descriptions of context models which need to be processed in a scalable fashion, and complex policy engines including conflict handling. For all these components, we have sketched starting points marking the current state-of-the-art.

Acknowledgment. The work presented in this paper has been funded by Science Foundation Ireland under Grant No. SFI/08/CE/I1380 (Lion-2) and SFI 08/SRC/I1403 (FAME).

REFERENCES

- [1] “Unified Communications Applications: Uses and Benefits,” Chadwick Martin Bailay, July 2008, http://www.cisco.mn/en/US/services/ps2961/ps2664/services_article_uc_apps_research_wp.pdf.

¹³Personal Information Model Ontology: <http://sourceforge.net/apps/trackoscaf/wiki/PIMO>

¹⁴The Nepomuk Semantic Desktop: <http://nepomuk.kde.org/>

¹⁵OASIS – Advancing Open Standards for the Information Society: <http://www.oasis-open.org/>

¹⁶<http://reverse.net/I2/software.html>

- [2] "Internet of Things – An action plan for Europe," Commission of the European Communities, June 18 2009, cOM(2009) 278 final.
- [3] J. Rosenberg, A. Houry, C. Smyth, and F. Audet, "Models for Intra-Domain Presence and Instant Messaging (IM) Bridging (draft-ietf-simple-intradomain-federation-04)," Internet Engineer Task Force, July 2009. [Online]. Available: <http://tools.ietf.org/html/draft-ietf-simple-intradomain-federation-04>
- [4] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson, H. Lu, X. Zheng, M. Musolesi, K. Fodor, and G.-S. Ahn, "The Rise of People-Centric Sensing," *IEEE Internet Computing*, vol. 12, no. 4, pp. 12–21, 2008.
- [5] H. Lu, W. Pan, N. D. Lane, T. Choudhury, and A. T. Campbell, "SoundSense: scalable sound sensing for people-centric applications on mobile phones," in *Proc. of the 7th International Conference on Mobile Systems, Applications, and Services (MobiSys 2008)*, K. Zielinski, A. Wolisz, J. Flinn, and A. LaMarca, Eds. ACM, 2009, pp. 165–178.
- [6] H. Junker, O. Amft, P. Lukowicz, and G. Tröster, "Gesture spotting with body-worn inertial sensors to detect user activities," *Pattern Recognition*, vol. 41, no. 6, pp. 2010–2024, 2008.
- [7] T. Stiefmeier, D. Roggen, G. Ogris, P. Lukowicz, and G. Tröster, "Wearable Activity Tracking in Car Manufacturing," *IEEE Pervasive Computing*, vol. 7, no. 2, pp. 42–50, 2008.
- [8] D. Choujaa and N. Dulay, "TRAcME: Temporal Activity Recognition Using Mobile Phone Data," in *2008 IEEE/IFIP International Conference on Embedded and Ubiquitous Computing*. IEEE Computer Society, 2008, pp. 119–126.
- [9] P. Mohan, V. N. Padmanabhan, and R. Ramjee, "Nericell: rich monitoring of road and traffic conditions using mobile smartphones," in *Proc. of the 6th International Conference on Embedded Networked Sensor Systems, SenSys 2007*, T. F. Abdelzaher, M. Martonosi, and A. Wolisz, Eds. ACM, 2008, pp. 323–336.
- [10] E. Miluzzo, N. D. Lane, S. B. Eisenman, and A. T. Campbell, "CenceMe - Injecting Sensing Presence into Social Networking Applications," in *Smart Sensing and Context, Second European Conference, EuroSSC 2007*, G. Kortuem, J. Finney, R. Lea, and V. Sundramoorthy, Eds., vol. 4793. Springer, 2007, pp. 1–28.
- [11] K. Aberer, M. Hauswirth, and A. Salehi, "Infrastructure for data processing in large-scale interconnected sensor networks," in *8th International Conference on Mobile Data Management (MDM 2007)*, C. Becker, C. S. Jensen, J. Su, and D. Nicklas, Eds. IEEE, 2007, pp. 198–205.
- [12] M. Sgroi, A. Wolisz, A. Sangiovanni-Vincentelli, and J. M. Rabaey, "A Service-Based Universal Application Interface for Ad-hoc Wireless Sensor Networks," in *Ambient Intelligence*, W. Weber, J. M. Rabaey, and E. Aarts, Eds. Springer, 2005, pp. 149–172.
- [13] J. Shneidman, P. Pietzuch, J. Ledlie, M. Roussopoulos, M. Seltzer, and M. Welsh, "Hourglass: An Infrastructure for Connecting Sensor Networks and Applications," Harvard University, Tech. Rep., 2004.
- [14] M. J. Franklin, S. R. Jeffery, S. Krishnamurthy, F. Reiss, S. Rizvi, E. W. 0002, O. Cooper, A. Edakkunni, and W. Hong, "Design Considerations for High Fan-In Systems: The HiFi Approach," in *CIDR 2005, Second Biennial Conference on Innovative Data Systems Research*, 2005, pp. 290–304.
- [15] P. B. Gibbons, B. Karp, Y. Ke, S. K. Nath, and S. Seshan, "IrisNet: An architecture for a Worldwide Sensor Web," *IEEE pervasive computing*, vol. 2, no. 4, pp. 22–33, 2003.
- [16] G. Mühl, L. Fiege, and P. Pietzuch, *Distributed Event-Based Systems*. Springer, 2006.
- [17] J. Rosenberg, "A Presence Event Package for the Session Initiation Protocol (SIP) (RFC 3856)," Internet Engineer Task Force, August 2004. [Online]. Available: <http://www.ietf.org/rfc/rfc3856.txt>
- [18] P. Millard, P. Saint-Andre, and R. Meijer, "XEP-0060: Publish-Subscribe," September 2008. [Online]. Available: <http://xmpp.org/extensions/xep-0060.html>
- [19] P. Bellavista, A. Corradi, and L. Foschini, "IMS-based presence service with enhanced scalability and guaranteed QoS for interdomain enterprise mobility," *IEEE Wireless Commun.*, vol. 16, no. 3, June 2009.
- [20] M. Altinel and M. J. Franklin, "Efficient Filtering of XML Documents for Selective Dissemination of Information," in *VLDB 2000, Proc. of 26th International Conference on Very Large Data Bases*, A. E. Abbadi, M. L. Brodie, S. Chakravarthy, U. Dayal, N. Kamel, G. Schlageter, and K.-Y. Whang, Eds. Morgan Kaufmann, 2000, pp. 53–64.
- [21] F. Fabret, H.-A. Jacobsen, F. Llirbat, J. Pereira, K. A. Ross, and D. Shasha, "Filtering Algorithms and Implementation for Very Fast Publish/Subscribe," in *SIGMOD 2001 Electronic Proceedings*, W. G. Aref, Ed., 2001, pp. 115–126.
- [22] G. Li, S. Hou, and H.-A. Jacobsen, "A Unified Approach to Routing, Covering and Merging in Publish/Subscribe Systems Based on Modified Binary Decision Diagrams," in *25th International Conference on Distributed Computing Systems (ICDCS 2005)*. IEEE Computer Society, 2005, pp. 447–457.
- [23] A. Crespo, O. Buyukkocuten, and H. Garcia-Molina, "Query Merging: Improving Query Subscription Processing in a Multicast Environment," *IEEE Trans. Knowl. Data Eng.*, vol. 15, pp. 174–191, 2003.
- [24] T. Gu, H. K. Pung, and D. Zhang, "A service-oriented middleware for building context-aware services," *J. Network and Computer Appl.*, vol. 28, pp. 1–18, 2005.
- [25] J. Coutaz, J. L. Crowley, S. Dobson, and D. Garlan, "Context is key," *Commun. ACM*, vol. 48, pp. 49–53, 2005.
- [26] J. Euzenat, J. Pierson, and F. Ramparany, "Dynamic context management for pervasive applications," *Knowledge Eng. Review*, vol. 23, pp. 21–49, 2008.
- [27] D. L. McGuinness and F. van Harmelen, "OWL Web Ontology Language Overview W3C Recommendation 10 February 2004," February 2004. [Online]. Available: <http://www.w3.org/TR/owl-features/>
- [28] N. Damianou, N. Dulay, E. Lupu, and M. Sloman, "The Ponder Policy Specification Language," in *Policies for Distributed Systems and Networks, International Workshop, POLICY 2001*, M. Sloman, J. Lobo, and E. Lupu, Eds., vol. 1995. Springer, 2001, pp. 18–38.
- [29] L. Kagal, T. W. Finin, and A. Joshi, "A Policy Language for a Pervasive Computing Environment," in *4th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2003)*. IEEE Computer Society, 2003, pp. 63–76.
- [30] A. Uszkowicz, J. M. Bradshaw, R. Jeffers, N. Suri, P. J. Hayes, M. R. Breedy, L. Bunch, M. Johnson, S. Kulkarni, and J. Lott, "KAoS Policy and Domain Services: Toward a Description-Logic Approach to Policy Representation, Deconfliction, and Enforcement," in *4th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2003)*. IEEE Computer Society, 2003, pp. 93–98.
- [31] T. Moses, "eXtensible Access Control Markup Language (XACML) Version 2.0," February 2005. [Online]. Available: http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf
- [32] V. Kolovski, J. A. Hendler, and B. Parsia, "Analyzing web access control policies," in *Proc. of the 16th International Conference on World Wide Web, WWW 2007*, C. L. Williamson, M. E. Zurko, P. F. Patel-Schneider, and P. J. Shenoy, Eds. ACM, 2007, pp. 677–686.
- [33] K. Barrett, "A Framework for the Semantic Translation of Policy Language Concepts," Ph.D. dissertation, Waterford Institute of Technology, 2009.
- [34] J. Strassner, J. N. Souza, D. Raymer, S. Samudrala, S. Davy, and K. Barrett, "The design of a novel context-aware policy model to support machine-based learning and reasoning," *Cluster Computing*, vol. 12, pp. 17–43, 2009.
- [35] A. Paschke, D. Hirtle, A. Ginsberg, P.-L. Patranjan, and F. McCabe, "RIF Use Cases and Requirements W3C Working Draft 18 December 2008," December 2008. [Online]. Available: <http://www.w3.org/TR/rif-ucr/>
- [36] J. de Bruijn, "RIF RDF and OWL Compatibility W3C Candidate Recommendation 1 October 2009," October 2009. [Online]. Available: <http://www.w3.org/TR/rif-rdf-owl/>
- [37] K. Verlaenen, B. D. Win, and W. Joosen, "Towards simplified specification of policies in different domains," in *Integrated Network Management, IM 2007. 10th IFIP/IEEE International Symposium on Integrated Network Management, Munich, Germany, 21-25 May 2007*. IEEE, 2007, pp. 20–29.
- [38] E. Lupu and M. Sloman, "Conflicts in Policy-Based Distributed Systems Management," *IEEE Trans. Software Eng.*, vol. 25, pp. 852–869, 1999.
- [39] S. Davy, B. Jennings, and J. Strassner, "The policy continuum-Policy authoring and conflict analysis," *Computer Communications*, vol. 31, pp. 2981–2995, 2008.
- [40] D. Lin, P. Rao, E. Bertino, and J. Lobo, "An approach to evaluate policy similarity," in *SACMAT 2007, 12th ACM Symposium on Access Control Models and Technologies*, V. Lotz and B. M. Thuraisingham, Eds. ACM, 2007, pp. 1–10.