SOCIETIES: Where Pervasive Meets Social

Kevin Doolin[1], Ioanna Roussaki[2], Mark Roddy[1], Nikos Kalatzis[2],
Elizabeth Papadopoulou[3], Nick Taylor[3], Nicolas Liampotis[2], David McKitterick[4],
Edel Jennings[1], and Pavlos Kosmides[2]

[1] TSSG, Waterford Institute of Technology, Waterford, Ireland
[2] National Technical University of Athens, Athens, Greece
[3] Heriot-Watt University, Edinburgh, UK
[4] INTEL Labs Europe, Leixlip, County Kildare, Ireland

Abstract. Traditionally, pervasive systems are designed with a focus on the individual, offering services that take advantage of their physical environment and provide a context-aware, personalised user experience. On the other hand, social computing is centred around the notion of a community, leveraging the information about the users and their social relationships, connecting them together often using different criteria that can range from a user's physical location and activity to personal interests and past experiences. The SOCIETIES Integrated Project attempts to bridge these different technologies in a unified platform allowing individuals to utilise pervasive services in a community sphere. SOCIETIES aims to use community driven context awareness, preference learning and privacy protection for intelligently connecting people, communities and things. Thus, the goal of SOCIETIES is to radically improve the utility of Future Internet services by combining the benefits of pervasive systems with these of social computing. This paper provides an overview of the vision, concepts, methodology, architecture and initial evaluation results towards the accomplishment of this goal.

Keywords: Pervasive Computing, Social Networking, Pervasive Communities, Cooperating Smart Spaces, Community Interaction Spaces, Future Internet.

1   Introduction

Pervasive computing [1] is the next generation paradigm in computer science that aims to assist users in their everyday tasks in a seamless unobtrusive manner, by transparently and ubiquitously embedding numerous computing, communication and sensing resources in the users' environment and devices. Until now, pervasive computing systems have been designed mainly to address the needs of individual users. This neglects an important part of human behaviour; socialising, and might partly explain the slow take-up of pervasiveness in commercial products. On the other hand, social computing [2] has enjoyed meteoric success in bringing people together online. Products in this area, however, do not integrate well with any but a few of the many devices and services to which their users have access.

This paper describes the work being carried out in the FP7 SOCIETIES (Self Orchestrating Community Ambient Intelligence Spaces) integrated project (www.ict-societies.eu), the aim of which is to investigate and address the gap between pervasive

and social computing by designing, implementing and evaluating an open scalable service architecture and platform for our so-called Pervasive Communities. Pervasive Communities have the potential to completely transform traditional online social networks, freeing them from web-applications and letting them loose in the real physical world. SOCIETIES supports the creation of purpose-driven Pervasive Communities by finding, connecting and organising relevant people and things from both physical and digital environments.

The core value proposition of Pervasive Communities is in continuous evaluation and refinement based on feedback collected from three real user trial groups that are involved in the project from the start and will be engaged until the project's completion. These user groups are: (i) the Student community, consisting of students from Heriot-Watt University in Edinburgh, (ii) the Disaster Management community of experts from the European Civil Protection Mechanism and (iii) the Enterprise community from Intel's offices in Ireland.

The purpose of this paper is to elaborate on the vision, concepts, methodology, architecture and initial evaluation results towards the realisation of Pervasive Communities. The rest of this paper is structured as follows: in Section 2, the SOCIETIES vision is presented and the concepts introduced are defined. Section 3 elaborates on the research challenges being investigated towards the Future Internet. In Section 4, the methodology adopted is described spanning from user research techniques, to technical requirement extraction methods and business analysis tools employed. Section 5 presents the architecture that has been designed and is being implemented, which exploits the benefits of both the pervasive computing and the social computing paradigms. In Section 6, the initial evaluation methodology and results are described. Finally, in Section 7, conclusions are drawn and future plans are exposed.

2 Vision and Concepts

2.1 The SOCIETIES Vision

While the majority of human social interaction takes place in the physical world, the digital world is becoming increasingly integrated into the social fabric in which we (co-)exist. The overarching goal of the EU funded project SOCIETIES is to seamlessly integrate the social aspects of our physical world with our digital equivalents.

The focus here is on the Discovery, Connection and Organisation of people and things (sensors, context data, devices, resources, services, information, etc.) into organised dynamically formed pervasive communities which provide their members with an enriched social experience, supported by enhanced proactive behaviour. The terms "discover", "connect" and "organise" describe the full functionality that supports the entire lifecycle of pervasive communities.

Discovery refers to the finding of people, communities, services, devices and resources across the physical and virtual worlds. What is critical to note here is that SOCIETIES provides the capability to discover entities that are relevant to a particular situation (be it a user's goals, desires, current context, etc.). This relevant discovery enables the specification of deep associations between a group of entities, with an external interface being made available in order to provide a third party

service provider, or simply a user, with access to that interconnected community. This pushes beyond the capabilities of current social networks and services which rely heavily on, for example, static personal information and user preferences, or manually provided context changes (such as a manual check-in). This allows for the provision of intelligent, rich, contextual data about users and the entities they interact with.
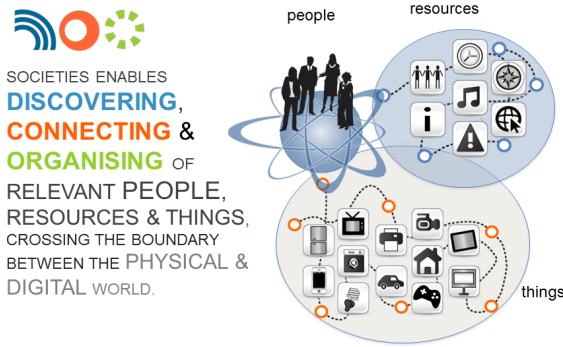

Fig. 1. The SOCIETIES vision

Harnessing these discovered entities, SOCIETIES then enables their connection across the physical and virtual worlds, allowing them to communicate and to capitalise on the capabilities of the entities each has formed a connection with as part of the pervasive community. This allows services such as crowdsensing, to take advantage of the multiple devices of the many members of a community and can provide them all with a much richer contextual picture of their physical and digital worlds than any one of them could generate alone. Trust, privacy and security are critical factors being considered in all phases of the project.

With connections in place, SOCIETIES can progress to organise and manage the lifecycle of this dynamically established community. This lifecycle management includes the introduction/removal of entities to/from the community – based on further discovery and connection cycles – and the formation of super-/sub-communities based on the context of the community as a singular entity. The results of the organisation allow for the creation of new service paradigms that use a community of connected entities as a focal point. The realization of pervasive community Discovery, Connection and Organisation in this fashion requires a significant effort in terms of research and development of new innovative functionalities, many of which provide value for third parties either as standalone or as more integrated intelligent features.

2.2   Concept Definitions

To address the vision above, the notion of Cooperating Smart Spaces (CSSs) has been introduced that aims to extend pervasive systems beyond the individual to dynamic communities of users. CSSs enable groups of users that demonstrate commonalities for a non-trivial period to join together in pervasive communities. To enable the merging of pervasive and social computing, in addition to the CSS notion, several new concepts have been introduced. These concepts are defined below.

A *Pervasive Community* is a group of, two or more, individuals who have agreed to share some, but not necessarily all, of their pervasive resources with other members of that community. The *Pervasive Resources* that can be shared are: (i) services, including services for controlling personal and environmental devices and (ii) information (both individual and community), including context, preferences, behaviours and memberships. A pervasive community, once constituted, forms a *Community Interaction Space* (CIS). There is a one-to-one mapping between pervasive communities and CISs. Individuals may belong to any number of pervasive communities, and thus CISs, simultaneously.

Members of a pervasive community interact with a CIS via their own personal *Cooperating Smart Space* (CSS). CSSs create the building blocks for enabling the integration of pervasive computing with social communities (physical or digital). CSSs constitute the bridge between a user's context (devices, sensors etc.) and the community the user is a part of. A CSS is a digital representation of a user or organisation, and also defines the impact that their services, information and resources have within a set of communities. As such, it represents the user's dynamic contribution to one or more communities. The CSS provides its owner with a suite of services which support the creation of, and participation in, pervasive communities as well as a range of intelligent cross-community functionalities, which enable the individual community member to benefit from the information and services of the community as a whole. A community is a collection of CSSs and/or supporting infrastructure services that wish to collaborate for mutually agreed purposes. There is a one-to-one mapping between individuals and CSSs. The only way in which an individual can participate in a CIS is via their CSS, but they can also interact with other CSSs without having to form pervasive communities or create CISs. Individuals may also interact with other individuals without using CSSs at all by employing more traditional mechanisms.

3   Research Challenges towards the Future Internet

SOCIETIES provides a platform that enables individuals to connect their physical activities with their online social environment in a transparent manner. Users avail themselves of the pervasive features of CSSs and CISs, such as context-aware, proactive and personalised service adaptation, implicit & explicit user behaviour learning, personalised privacy protection and cross domain access to services & resources. On an individual basis, pervasive information, such as user behaviour models and context information, are enhanced with knowledge inferred by monitoring the social interactions of the user on the social networks they visit. Augmenting the personalisation and context information improves the quality of the pervasive technologies that CSSs offer. On a community level, SOCIETIES utilises data such as user behaviour, context and trust to form communities of users who share similar interests, hobbies, careers, etc., enabling free exchange of information and knowledge. Moreover, it provides an open scalable service architecture that allows users to share resources (e.g., services, devices, sensors) with members of the same community equipped with a range of resource sharing policies for resolving conflicts. Furthermore, services and resources can be configured at both CSS and CIS level by taking into account community context and behaviour that is inferred by collecting and mining information from the members of the community.

The most essential resource for realising the CSS vision is the availability of information and more importantly personal information. To address the issue of privacy that rises, CSSs provide a range of intelligent privacy protection techniques for managing the flow of information and allowing the user to have complete control over the handling and disclosure of their information. Users can explicitly create privacy preferences that state how their information is disclosed while the system is also able to implicitly learn privacy preferences by monitoring the user's behaviour related to privacy protection.

Clearly SOCIETIES draws together a number of key challenges for the Future Internet. Social computing, in many different contexts and through various devices, is becoming a major driver for Internet use. Pervasive systems, as embodied in smart spaces, are also set for a deployment explosion and will capitalise on the Internet of Things to make enormous demands on the Internet in the near future. The social and the pervasive aspects of Internet use each raise important privacy challenges in their own right but together, the risks and consequences of failing to provide adequate and usable privacy mechanisms increase exponentially. In addressing all of these challenges SOCIETIES is making a significant step towards shaping the Future Internet.

4   Methodology

4.1   User Research Methodology

The user research methodology adopted was required to meet the following objectives: (i) introduce users to CSS concepts and novel technologies, (ii) motivate users to participate through engaging with manifestations of CSS visions that are meaningful to them, (iii) yield requirements from user research that would be initially informative to project development, (iv) enable researchers to have some useful and empathetic insight into the lives, interests, and concerns of potential users, (v) utilise research activities that could be conducted in a minimal amount of time with little effort or disruption to the lives of the users and (vi) facilitate equal access interaction between the stakeholders of user groups and researchers with malleable scenarios.

The user research methodology employed a triangulation of methods (Figure 2) that included: observation via ethnographic methods, self-reporting via online surveys and scenario led participatory workshops. For each group, the approaches used varied taking into consideration the location, environment, user access and availability. Results from the three different approaches of user research were presented as ethnographic vignettes, statistical analysis and updated scenarios. User requirements were extracted from combining these three sets of results for each of the three user groups.

Rapid ethnographic techniques, such as participant observation, contextual enquiry, shadowing and guerilla fieldwork were employed to observe people in their natural environment, thus gaining insights and understanding about the everyday worlds of users in each group. The results of this research were presented as ethnographic vignettes and field observations.

An *online questionnaire* was formulated to gather information about each user group's general demographics; their knowledge and experience of technical tools, pervasive features, and social networking services; their current community connections, and other relevant information specific to each group. The survey was
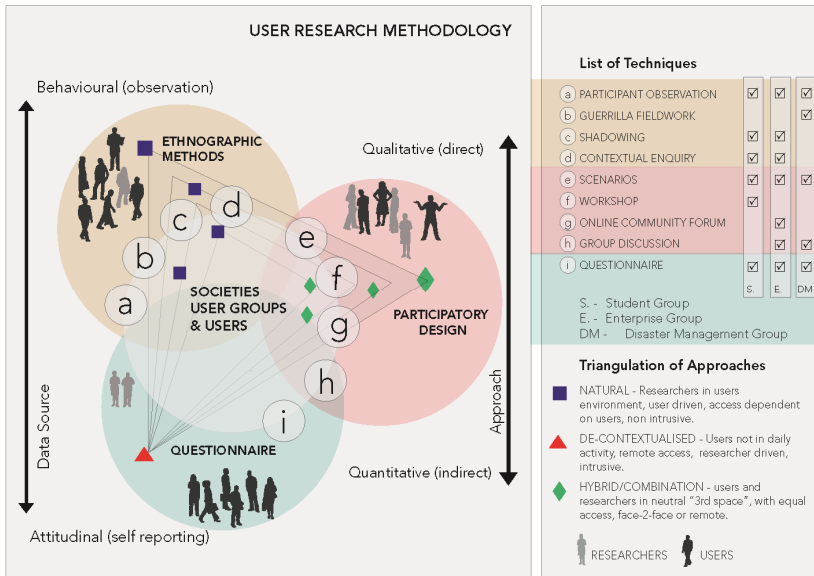
Fig. 2. User Research Methodology

sent to a random sample from each group. The responses were anonymised. Graphs of the results demonstrate commonalities and differences in the social use of technology within and across the groups.

Participatory Design (PD) ***workshops*** were organized to provide a democratic, collaborative approach, facilitating creative, cooperative involvement of all the stake-holders in the development of project concepts and services. Scenarios were selected as a key tool for the PD sessions, as they function both, as a creative process for visioning exercises, as well as an empathetic narrative conduit for complex ideas and information. Initial scenarios demonstrating possible uses of the proposed platform in the context of student, enterprise and disaster management situations were sketched in brainstorming sessions with researchers. These initial scenarios were in turn introduced to users by researchers in the neutral creative third space of PD workshops [3], where participants' reactions, ideas and discussions led to alterations and advancements of these scenarios. Creative understandings [4] forged in these sessions led to updated scenarios envisioning how pervasive communities could function in each group's social setting.

4.2   Technical Requirement Extraction Methodology

Based on an evaluation of the state of the art methodology approaches, it has been decided to classify the technical requirements in a manner similar to the approach suggested by the FURPS model [5]. More specifically, a scenario-driven process has been followed to collect and specify the technical requirements. The five stages of this elicitation process are illustrated in Figure 3 and are briefly described below.
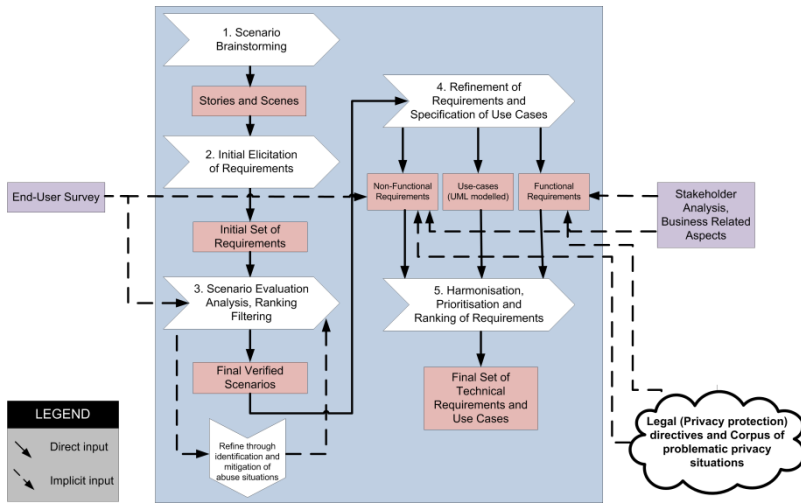
Fig. 3. Technical Requirements specification methodology

1. ***Scenario brainstorming:*** This stage aimed to the design of various story flows and scenes demonstrating and extending the features of the envisaged system.
2. ***Gathering of initial requirements:*** In this stage, an initial set of functional requirements was extracted from the scenarios produced in stage 1. These requirements were classified in five main categories: General, Deployment, Service and Resource, User Experience and Security-related Requirements.
3. ***Scenario evaluation, analysis, ranking, filtering and refinement:*** In this stage, the initial scenarios produced in stage 1 were evaluated and ranked based on various criteria, such as the volume of features they demonstrated compared to the feature set captured in the vision of Section 2, the quality of the initial requirements collected in the second stage, etc. Based on this evaluation, on the end-user feedback collected (§ 4.1) and on the business analysis performed (§ 4.3), a set of refined final scenarios was produced.
4. ***Refinement of functional and non-functional requirements and extraction of use cases:*** In this stage, the final scenarios have been studied in order to extract additional technical requirements (both functional & non-functional), as well as use-cases, while the initial requirements collected in stage 2 were homogenised, merged, eliminated, extended and classified.
5. ***Harmonisation, prioritisation and ranking of requirements:*** In the final stage, the elicited requirements were prioritized, harmonised and checked for consistency.

4.3   Business Analysis Methodology

The business analysis process took place in two phases. First, an adjustment of a subset of the Tropos methodology [6] was used to extract business opportunities and potential revenue streams. This phase included five distinct steps: (i) market analysis that identified the existing stakeholders and business models related to the investigated domains, (ii) scenario analysis that processed the final scenarios

produced for the technical requirements extraction in order to identify the related stakeholders, the potential business interests that arise and the business opportunities that emerge, (iii) generalization of the identified stakeholders that led to the identification of the existing and new stakeholders that are involved in the envisaged system, (iv) extraction of business requirements that are stakeholder specific and (v) extraction of the business opportunities and the respective value proposition. Once the process above was complete, the Business Model Canvas methodology [7] was exploited to assist in defining the applicable business models. Thus, the business model canvas approach has been used over the Discover, Connect and Organise phases that contribute to the formation of the envisaged system. This resulted in identifying how this system can offer value to various stakeholders; portray the capabilities and partners required for creating, marketing, and delivering this value, with the goal of generating profitable and sustainable revenue streams.

5 Architecture

The architecture that implements the concepts presented above is illustrated in Figure 4, where an overview of the "core services" provided by the proposed architecture is provided. The services depicted are grouped according to the major concept they manipulate or operate on. Thus, services that operate on a single CIS are grouped together, as are those that operate on a CSS, and those found on every node in a CSS.
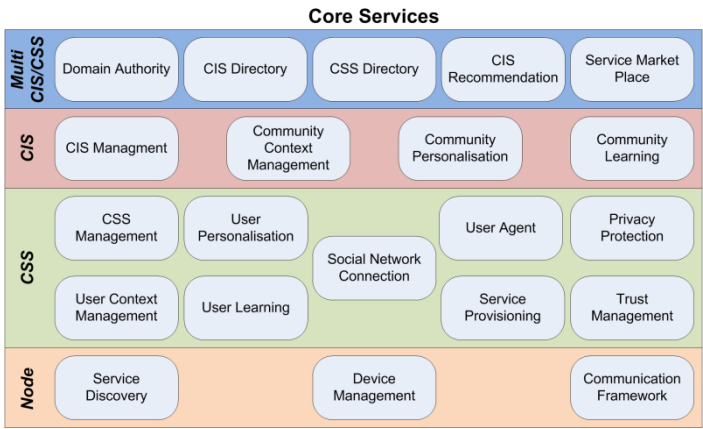
Fig. 4. The Cooperating Smart Space layered functional architecture

Multi CSS/CIS Services operate for the benefit of more than one CIS, or more than one CSS. Thus, they operate for a wider group of stakeholders. They offer federated search and domain administration functions and require multiple CSSs or CISs to be effective. This group includes the following services: the *Domain Authority* (that provides and manages the CSS and CIS identities in a decentralised manner, allowing authentication between multiple domains), the *CIS Directory* (that manages the CIS information in a decentralised repository, it records available CISs within a domain or set of domains, it enables searching for CISs based on specific criteria and it allows a CIS to be removed from the repository), the *CSS Directory* (that provides search

facilities for CSSs, based on their identifier or by specifying search criteria, such as, public profile attributes and tags), the *CIS Recommendation* (that is responsible for handling CIS recommendations, allowing for recommendations of CISs to users and vice versa, considering, among others, the users' privacy preferences) and the *Service Market Place* (that provides access to a repository of installable 3^{rd} party (3P) services and optional "core" services and provide mechanisms for accounting and charging).

   CIS Services operate on behalf of a single CIS. There is at least one instance of these services per CIS and an instance of these services can be used by multiple CISs. The CIS services are: the *CIS Management* (that is responsible for handling all aspects of CIS lifecycle management (creation, update and removal), provides control over CIS membership and includes a community profile manager and a role manager to specify the governance model for the CIS), the *Community Context Management* (that enables access to and maintenance of community context, providing query capabilities, as well as, addition/update/removal operations for community context, maintaining the history of context for a CIS, and inferring community context information), the *Community Learning* (that supports community preferences and community intent learning) and the *Community Personalisation* (that manages the community preferences and community intent and exposes interfaces for community members to retrieve these preferences and intent models for their own use).

   CSS Services operate on behalf of a single participant or CSS. There is at least one instance of these services per participant and an instance of these services can be used by multiple participants. The CSS services are: the *CSS Management* (that controls which Nodes (devices or cloud instances) are part of the CSS, assigns a common identifier and manages resource sharing & configuration policies), the *User Context Management* (that is responsible for acquiring the user context from sensors and other context sources, for modelling and managing the collected data, for maintaining current & historic context in appropriate data repositories and for the provision of inference techniques enabling the extraction of high level information from raw context data), the *User Personalisation* (that manages & evaluates the user behavioural models, such as user preferences, user intent, Bayesian models, etc., and eventually identifies the actions that need to be taken), the *Social Network Connection* (that integrates with existing Social Networking Systems (SNSs), enabling the extraction of public info available in SNSs, as well as access/update of non-public information for the specified user), the *Privacy Protection* (that provides identity management mechanisms, facilities for managing the CSS privacy policies, which specify the terms and conditions the CSS will respect concerning the personal data, also offering Privacy Policy Negotiation facilities), the *User Learning* (that supports learning of user behaviour models exploiting the user's history of actions stored in the system), the *User Agent* (that acts on behalf of a single CSS based on information from several CSS and CIS components, aiming to establish the system's proactive behavior, resolving any conflicts that may arise, also enabling CSS users to provide feedback on the system actions or decisions), the *Trust Management* (that is responsible for collecting, maintaining and managing all information required for assessing the trust relationships and includes a Trust Engine for evaluating direct, indirect and user perceived trust) and the *Service Provisioning* (that supports the setup and lifecycle control of a 3P service or CSS resource, allowing for installation, (re)configuration and removal of new 3P services, also supporting the enforcement of 3P service sharing policies).

Node Services are available per CSS Node. A CSS Node is a logical node (device or cloud instance) running CSS software that coordinates with other CSS Nodes to form a participant's CSS. There is an instance of these services per CSS Node. This grouping includes the following services: the *Communication Framework* (that provides the necessary mechanisms to support intra- and inter-CSS communication, supporting the identification and maintenance of network connections, the discovery of CSS Nodes (devices), and the communication between discovered nodes), the *Device Management* (that provides mechanisms for managing devices within a CSS, supporting the discovery of hardware devices and management of their capabilities) and the *Service Discovery* (that provides service discovery and advertisement mechanisms, enabling the discovery of core platform services within a CSS, as well as, the discovery of 3P services shared by other CSSs or CISs).

6   Initial Evaluation

Using Paper Trials, an initial user evaluation was conducted in April 2011 across all three user communities, i.e., the Disaster Management, the Student and the Enterprise community. The primary objective of these trials was to record users' responses to early prototypes of initial scenarios & concepts and how users' experiences of these prototypes conformed to the previously identified user requirements. These Paper Trials were interpreted loosely as a user evaluation trial of low-fidelity prototypes. A secondary objective was to engage with users to confirm or discover the opportunity spaces for pervasive and social computing, "*where there is no urgent problem to be solved, but much potential to augment and enhance practice in new ways*" [8].

The envisaged system posed a challenge that could not be served by traditional paper prototyping alone, since it required prototypes for user evaluation that focused on user activities, goals and contexts of use, with varied levels of detail, thus conveying a range of CSS/CIS system interactions within the user domains, which were not necessarily focused on users manipulating device interfaces (i.e., pervasive services working in the background for the benefit of their users). Therefore a specific evaluation methodology was necessary.

6.1   Evaluation Methodology

Two user evaluation methods were selected for the trials, namely Storyboards and Wizard of Oz. Both methods are known to be well suited to work with the scenario-based design approach method, which was employed to describe potential deployment for each of the three user communities. Both are methods that employ scenario based vision prototypes [9], which serve the purpose of defining early design focus for developers and providing a site for evaluating user responses [10].

Storyboards were the primary method used and acted as an informal low-fidelity design artefact that provided a concrete but flexible example of how services and features of the envisaged system, in this case an intelligent pervasive communities system deployed in a context familiar to each user group, would be envisioned by the consortium. The users' feedback to our storyboards was captured using two different techniques: (i) questions that were embedded in each storyboard and were addressed to individual users in the context of viewing the storyboard, to focus their attention

and allow them to directly answer questions of particular interest to the project researchers and (ii) participatory discussions that were facilitated after the storyboard viewings, in the case of the Enterprise and Disaster Management communities, where the storyboards provided springboards to openly discuss users' reactions to the issues and scenarios depicted within.

The *Wizard of Oz* method was used as a secondary method in the case of the Student community. This method utilized a script based on university scenario segments, with the project researchers playing the role of the envisioned system by managing environmental and device responses to user activities and preferences in a pervasive laboratory, which had been set up to stage an intelligent campus environment. Participants answered questions posed during the experiment that was also videotaped. It was designed to allow students evaluate an immersive experience of a social and pervasive environment.

6.2   Evaluation Results

The overall feedback from this evaluation study did indicate strong support for the concepts that were presented, albeit with quite a number of concerns expressed, including: trusting the system, controlling privacy and difficulties with accepting automation. It was evident that users did see the value concept of creating purpose-driven communities and the leveraging of collective intelligence from those communities, which is at the core of the project's value proposition.

Regarding privacy and information disclosure, most students were happy to disclose basic information. Most users were in favour of sharing their preferences and requirements, although they were more reluctant to disclose other information and wanted to be consulted on disclosures rather than letting their devices make decisions for them. This was part of a general concern regarding trusting a system to make automatic decisions on behalf of users.

There was a common concern between the student and enterprise users on how far technology should go in replacing natural human behaviours. The students considered it acceptable if the system suggests greeting somebody based on shared interests or intents, as well as the person's mood, but only a few students would use this to start friendships, preferring to meet new people "the old-fashioned way" instead. The enterprise users liked the professional networking support features, but some questioned the effect it may have on the natural networking activity and suggested it could rule out opportunistic encounters. One user stated: "*A user could become a slave to their preferences and may lose the opportunity to discover new opportunities*".

It is clear that a main objective for the project's researchers is to design a system for creating purpose-driven communities, through rich context data sets. Yet, it is also clear from our user responses that a key concern is that the system appears to be too intrusive and that there would be serious concerns around the areas of privacy, trust, automated community creation and service delivery. This appears to be a paradox, since the system requires the user to give rich context data in order to provide context-aware personalised services and as a result is a key challenge for the consortium.

7   Conclusions

The SOCIETIES project aims to investigate and address the gap between pervasive and social computing by designing, implementing and evaluating an open scalable

service architecture and platform. Based on a vision of the discovery, connection and organisation of relevant people, resources and things into dynamically formed pervasive communities, SOCIETIES attempts to bridge the domains of pervasive and social computing in a unified platform allowing individuals to utilise pervasive services in a community sphere.

This paper presents concepts and research methodologies adopted in the SOCIETIES project towards the realization of Pervasive Communities and in order to assess whether real end users can see a value for engaging with such a system. The overall feedback from our initial user evaluation study did indicate strong support for the concepts that were presented, albeit with quite a number of concerns expressed, including: trusting the system, controlling privacy and accepting automated decision making. Based on these user concerns and considering every user feedback collected from the initial trials, the technical requirements have been revised and the SOCIETIES architecture has been adapted accordingly. Two more user trials have been scheduled for 2012 and 2013 that will enable us to assess how successfully the vision, concepts and results of SOCIETIES address the technological and user acceptance gap between pervasive and social computing. The results achieved up to this point and the user feedback already collected indicate that the SOCIETIES platform can find its way to the facilities portfolio that users exploit on a daily basis.

References

1. Saha, D., Mukherjee, A.: Pervasive computing: A paradigm for the 21st century. IEEE Computer 36(3), 25–31 (2003)
2. Dasgupta, S.: Social Computing: Concepts, Methodologies, Tools, and Applications. IGI Global (2010)
3. Muller, M.J.: Participatory Design: The third space in HCI. In: The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications, pp. 1051–1068. Lawrence Erlbaum Associates, New Jersey (2002)
4. Wright, P., McCarthy, J.: Empathy and Experience in HCI. In: 26th Annual SIGCHI Conference on Human Factors in Computing Systems, Florence, Italy (April 2008)
5. Grady, R.B., Caswell, D.L.: Software Metrics: Establishing a Company-Wide Program. Prentice Hall, Englewood Cliffs (1987)
6. Castro, J., Kolp, M., Mylopoulos, J.: Towards Requirements-Driven Information Systems Engineering: The Tropos Project. Information Systems 27(6), 365–389 (2002)
7. Osterwalder, A., Pigneur, Y.: Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers. John Wiley & Sons, New Jersey (2010)
8. Hornecker, E., Halloran, J., Fitzpatrick, G., Weal, M., Millard, D., Michaelides, D., Cruickshank, D., Roure, D.D.: UbiComp in opportunity spaces: challenges for participatory design. In: 9th Conference on Participatory Design, Trento, Italy (August 2006)
9. Erickson, T.: Notes on Design Practice: Stories and Prototypes as Catalysts for Communication. In: Scenario-Based Design: Envisioning Work and Technology in System Development. John Wiley & Sons, New York (1995)
10. Houde, S., Hill, C.: What do prototypes prototype? In: Handbook of Human-Computer Interaction, 2nd edn., Elsevier Science B. V., Amsterdam (1997)