Chapter 7
Enhancing Mobile Social Networks with Ambient Intelligence

Kevin Doolin, Nick Taylor, Micheal Crotty, Mark Roddy, Edel Jennings, Ioanna Roussaki, and David McKitterick

Abstract Mobile social computing has exploded into people's lives during the past 10 years, but to become truly pervasive it needs to be much more context-aware and personalizable. The next generation of social media needs to be able to react and adapt to the physical environments in which people live and act. The SOCIETIES project is integrating research undertaken in the field of pervasive computing with social computing to develop the next generation of social media systems. The vehicle for this is the "pervasive community", and this chapter outlines the innovations required to realize this concept. Pervasive communities can restore the symbiosis between our digital and physical worlds.

7.1 Introduction

When considering the future of mobile social computing, it is important to recall that people are mobile creatures, and from the time we developed our first indirect means of communicating (writing) the tools we use to socialize have

K. Doolin (✉) • M. Crotty • M. Roddy • E. Jennings
Telecommunications Software and Systems Group (TSSG), Waterford Institute of Technology, Waterford, Ireland
e-mail: kdoolin@tssg.org; mcrotty@tssg.org; mroddy@tssg.org; ejennings@tssg.org

N. Taylor
School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, UK
e-mail: N.K.Taylor@hw.ac.uk

I. Roussaki
National Technical University of Athens, 157-173 Zographou, Athens, Greece
e-mail: ioanna.roussaki@cn.ntua.gr

D. McKitterick
INTEL, Collinstown Industrial Park, Leixlip, County Kildare, Ireland
e-mail: davidx.mckitterick@intel.com

A. Chin and D. Zhang (eds.), *Mobile Social Networking: An Innovative Approach*, Computational Social Sciences, DOI 10.1007/978-1-4614-8579-7_7,

accommodated our roaming nature. Historically, those interludes during which humanity has surrendered to attending fixed locations in order to socialize indirectly have been brief. The fixed telephone was supplemented by the mobile phone within 50 years of it entering most homes. The early years of social networking, which tied people to their desktop computers, were just another such interlude. The users of those sites already had SMS available to them when they were on the move, and mobile social networking was inevitable with the advent of the smartphone.

Pervasive computing has the potential to totally transform traditional Web-based social networking sites (SNSs), freeing them from the Web-browser and letting them loose in the real world (Pallis et al. 2011; Xu et al. 2011). The most popular of these, such as Facebook, LinkedIn, MySpace, and Twitter, have already embraced smartphone technology, offering mobile versions for their users to access while on the move (Semertzidis and Laso-Ballesteros 2010). Future social networking users will demand the context-aware "everyware" services supporting personalization, which pervasive computing can provide anywhere and at any time.

The SOCIETIES project (SOCIETIES 2012) aims to bridge the gap (Doolin et al. 2012) between the current Web and smart phone application-based social networking services (Zaphiris and Ang 2009) and the next generation of pervasive computing services (Obaidat et al. 2011; Hansmann et al. 2012). In its simplest manifestation, the integration of pervasive and social computing will permit the online "poke" or "nudge" to morph into an image of the poker appearing on a display in the poked person's vicinity, wherever they are. This might still be a computer screen but it could also be a public display, TV, digital photoframe, a haptic or audio notification on a mobile device, etc. More advanced convergence will unify the two concepts completely into "pervasive communities". Mobile devices form a core element of this pervasive community concept. It is envisaged that the mobile devices of a user will be fully integrated into a community of common mobile resources, applications, and services – in essence, the mobile device will provide the user's hook into the pervasive community, will provide their identity, and will provide aspects of their context. As such, the community will always follow the user (based on their current context), and will integrate with fixed and mobile smart spaces as required. This fully ties into the concept of local–social–mobile.

The people in one's social network could be locatable and, if in sight, identifiable, within the limits of the personal privacy policies of the individuals concerned, of course. People with shared interests are likely to belong to a common pervasive community or group within an SNS. They are likely to attend the same events or places of interest and, if they were at such a place concurrently, pervasive communities could alert them to each other's presence and enable them to meet physically. Pervasive communities will enable people to leave messages about exhibits, restaurants, etc. which can be picked up automatically by other members who visit the same place at a later date. Streaming video facilitated by pervasive communities will enable their members to see and hear what their friends are seeing and hearing live with no effort and no delay.

The context-awareness of pervasive communities is not simply limited to location-awareness, however (Lukowicz et al. 2012). The "everyware" (Greenfield 2006) information available to a pervasive community can report a host of measurable characteristics about a member's environment and pass them to other members, even when on the move. As previously stated, it is the mobile device that is the hook that associates the user to their pervasive community. Weather conditions, for instance, helping one to decide where to go and what to wear on a free afternoon, as well as traffic reports, queue lengths at tourist attractions, delay times for planes, trains, and buses, and even dynamic information of other community members, such as their current activity, status, and location, are all types of information that can be made available to a pervasive community member anywhere anytime, either in their raw form, or filtered through the perceptions of what like-minded fellow members believe to be important and/or relevant.

Nor are the actions which a pervasive community member can take limited to data manipulation. Pervasive services will enable a member to affect things in their environment and, if one was minded to, one could allow one's friends to change one's environment using those same services. A friend could change one's TV channel to a particular show that they have suggested you watch, for instance or set one's oven to a temperature just right for baking the perfect cake. Robotic companions could be controlled by friends who are geographically separated. Computer-supported co-operative work (CSCW) could be completely transformed by the unification of pervasive and social computing in pervasive communities.

The rest of this chapter is structured as follows. Section 7.2 elaborates on the vision of the SOCIETIES project introducing the *discover–connect–organize* approach. Section 7.3 defines the necessary concepts and presents the SOCIETIES system architecture. Section 7.4 describes the main innovations of the SOCIETIES project, classifying these into seven categories. Section 7.5 elaborates on the SOCIETIES methodology that is based on the inclusion of three different user communities: an Enterprise, a Student and a Disaster Management Community. Section 7.6 provides some details on the project status, as well as on the future research and exploitation plans. Finally, in Sect. 7.7, conclusions are drawn.

7.2 SOCIETIES Vision

To date, pervasive computing systems have been designed mainly to address the needs of individual users. The sensor nets and services, which form part of the smart spaces of these pervasive systems, while supporting multi-user operations, are targeted at people who are assumed to be only interacting with the smart space and not directly with other users. At the same time, the recent development and rise in popularity of social computing has occurred more or less in isolation from developments in pervasive computing (Parthasarathy et al. 2011).

Experience has taught us that some of the most beneficial and rapidly assimilated uses of computer technology have arisen when that technology has brought people together and allowed them to communicate and collaborate. Computer-supported co-operative work (CSCW), massively multiplayer online role-playing games (MMORPGs), social networking sites (SNSs), even the humble SMS or instant messenger service have seen meteoric rises in popularity because they address a deep-rooted human need to socialize. The integration of pervasive technologies, from the real physical world, with social computing has the potential to generate a new wave of innovations in this novel hybrid area, and to re-factor ICT systems at the point of use so that they become more user-friendly, responsive, relevant, and capable of providing new means to harness data and engage more people, with a view towards tackling societal problems.

The SOCIETIES project aims to address the gap between pervasive and social computing by designing, implementing, and evaluating an open scalable service architecture and platform for pervasive communities. A pervasive community is inherently context-aware, and so can adapt to factors such as the user's location, activity, environment, physiological/psychological status, and current goals. In order to achieve this, a pervasive community has to be self-organizing, self-improving and capable of proactive behavior, which will optimize and personalize the pervasive experience of an individual user or an entire community.

A growing challenge that is faced by both social and pervasive computing is the enormous volume and variety of resources available to their users. As more people participate in social media, and as those media provide richer ways in which people can interact and share content, it has become essential to provide mechanisms which enable users to filter and organize the content they receive, in order to home in on what is relevant to them. Similarly, as pervasive technologies make more, and richer, sensors, networks, and services available to people at any time and in any place, identifying what is relevant, and what is not, becomes a fundamental requirement for usability.

The SOCIETIES platform recognizes, the paramount importance of relevance and the ways in which social and pervasive information and methods can be brought together to provide mutual support in identifying what is relevant to a user at any given time. Examples of factors that can contribute to determining what is relevant at a time, *t*, are:

- *Location* – at current time (t), historically at (t−1, t−2, ..), predicted at (t+1, t+2, ..)
- *Activity* and *status* – recognized by physical action monitoring and digital service usage, or predicted from prior tasks at (t−1, t−2, …)
- *Interactions* – people with whom one is communicating at present (t) or in the recent past (t−1, t−2, …), people with whom one is co-located, people with whom one is sharing services
- *Profile* – interests, background, affiliations, etc.
- *Physiology* – body temperature, perspiration rate, blood sugar level, etc.
- *Psychological/emotional state* – (un)happy, angry, bored, interested, etc.
- *Preferences* – service, interface, group preferences, etc.

Evidence from such a wide variety of sources (some pervasive, some social) can be bound by confidence limits, and combining the degrees of confidence from all sources can permit the calculation of an overall "quality of relevance" (QoR). This QoR can, in turn, be used by pervasive and social computing systems to decide what to filter out and what to alert a user to at a specific time and/or place. Such a decision process also needs to be personalizable, of course, so that each user can focus it to his or her own needs and tastes in a context-dependent fashion.

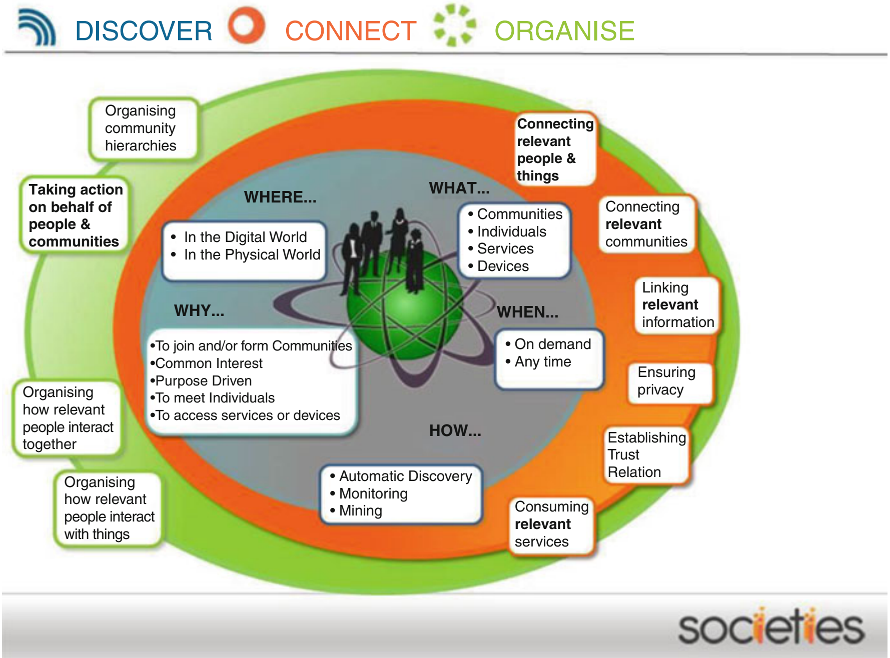

Fig. 7.1 The discover, connect, organize approach of SOCIETIES

The vision of SOCIETIES arises from this requirement to distinguish the relevant from the irrelevant, and can be summarized in three key concepts, each of which contributes to the formation of our pervasive communities. These concepts are: *discover*, *connect* and, *organize*, and the respective approach is illustrated in Fig. 7.1. In short, *SOCIETIES enables the discovery, connecting, and organizing of relevant people, resources, and things, crossing the boundary between the real and virtual worlds*. They are elaborated in turn in the subsequent sections.

7.2.1 *Discover*

SOCIETIES enables the discovery of trusted entities that reside in both the physical and digital worlds, such as individual people, communities, devices, resources, and services. What is important to note here is that the discovered entities have a distinct relevance either to each other or to the user on whose behalf the discovery process is operating – this concept of relevance is key throughout all innovations realized by the project, and sets SOCIETIES apart from current systems which typically adopt an all or nothing approach. Determining relevance is elaborated on further in the innovation sections that follow.

Existing systems such as LinkedIn or Facebook facilitate the discovery of people based on lightweight personal preferences and associations (primarily through manipulation of friends lists), while others harness lightweight contextual information (for example the use of a manual check-in on a mobile device); however, these systems do not enable the full exploitation of the capabilities of pervasive technologies or sensor networks to discover entities across the virtual – physical worlds of a user. SOCIETIES enables, for example, the discovery of a community of people with a common interest without a dependence on social network information (although it has to be noted that SOCIETIES integrates information from social networks in order to provide a deeper association between entities). Similarly, SOCIETIES can enrich existing social network services through the provision of rich contextual data about their users. The discovery notion in SOCIETIES therefore goes beyond the capabilities of existing technologies. Goal-, performance-, intention-, and commonality-driven discovery can greatly increase the performance of systems where access to relevant but unknown entities is required.

In terms of convenience, SOCIETIES' capability to discover based on relevance will increase the level of convenience provided to the entity seeking connections.

The discovery system is highly personalizable and context-aware. It can provide goal/social connection-, or situation-based discoveries, as well as learning and taking proactive discovery actions on behalf of users. Discovery also takes users' privacy requirements into consideration as SOCIETIES adopts a "privacy-by-design" approach throughout. Through the use of open federated identity management services, SOCIETIES allows the current social network monopoly to be broken, so that users can decide who provides their identity service. This will improve the overall quality of service and promote diversity of business models and opportunities.

Note that the provision of a privacy-aware external interface to allow third-party service developers to access this rich set of relevant entities opens up new possibilities in terms of next-generation mobile and social computing applications.

7.2.2 Connect

Building on the discovered relevant entities, SOCIETIES provides mechanisms for the interconnection of these entities across the physical and digital worlds. This enables communication connections to be established which capitalize on the capabilities of the entities they connect to. From a user's perspective, connections can take many forms – person to person, person to group, person to object, person to service – and only by accommodating all of these can any individual user's digital and physical worlds be seamlessly bridged. With the provision of external interfaces, third-party service providers are able to use SOCIETIES to enter the community-based mobile services market. Additionally, management of community resources, and off-device aggregated analytics, will enhance system and device performance (Lukowicz et al. 2012).

Convenience for the end user is greatly enhanced, since SOCIETIES enables the pervasive notion of having a user's/community's resources interconnected seamlessly, which alleviates the need for complex manual setup of entities.

Of course, the interconnections can be personalized and made context-aware on an individual basis, as will be discussed in the innovations section later.

7.2.3 Organize

The final stage of the DCO concept is of course "organize" (note that although the DCO concepts are presented in a linear manner in this document, in reality they form part of a continuous cycle).

Organizing, from a SOCIETIES perspective, relates to the complete lifecycle of discovered and connected entities. This lifecycle management includes the introduction of new entities to a connected community (based on further discover and connect cycles), and removal of entities that are no longer relevant or which may no longer require to be part of a specific community.

Enhanced intelligence functionality is available to community entities, and can be enriched by third-party services that can leverage SOCIETIES' enabling technologies such as privacy/trust management, security, learning, personalization, and context-awareness. It is critical to note here that once a community is formed, SOCIETIES then enables the deployment of new community services that harness contextual information from the physical and digital world, which can be fully personalized on an individual and community basis, and which protect privacy at all times. Some simple examples include the organization of a dynamic community for a business meeting or conference, or organizing a group of experts when needed in a disaster situation.

SOCIETIES' support for community hierarchies, members, and lifecycles facilitates the formation of temporary and ongoing communities for example. Analysis of an ad hoc temporary community's activities can facilitate the creation of a more permanent/ongoing community (for example a temporary community formed in a museum could become an ongoing community interested in a particular era of history). Additionally, identification of cohesive sub-communities can lead to the formation of useful community hierarchies (for example, a community of students could be formed into sub-communities based on course content, and therefore receive more focused, relevant services based on learning materials).

7.3 SOCIETIES Framework and System Architecture

7.3.1 Core Concepts

The project is formed around three interdependent concepts, illustrated in Fig. 7.2.

Pervasive community – a group of two or more individuals who have agreed to share some, but not necessarily all, of their pervasive resources.

CIS – a pervasive community, once constituted, forms a community interaction space (CIS). Individuals may belong to any number of pervasive communities and thus CISs simultaneously.

CSS – members of a pervasive community interact with a CIS via their own personal cooperating smart space (CSS), which lives on their mobile device(s). They can also interact with other CSSs directly, or without using CSSs at all. People can interact in person.

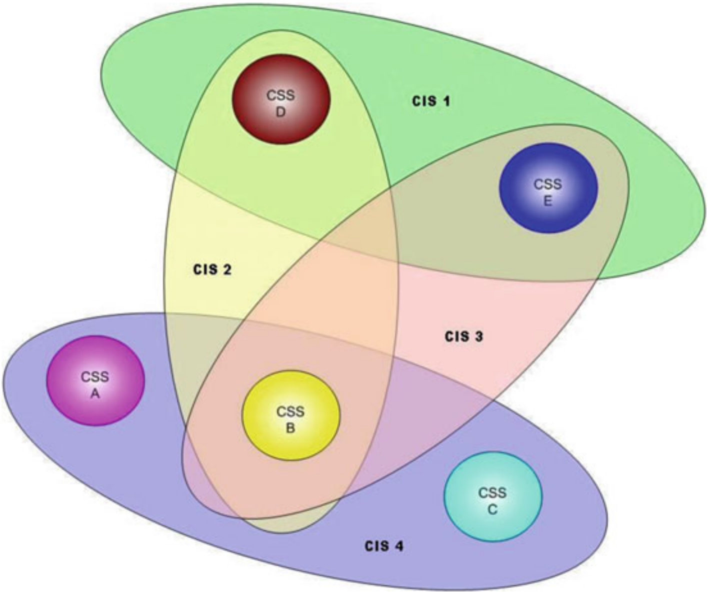

Fig. 7.2 Five individuals using their CSSs to form four pervasive communities

7.3.2 Architecture

This section provides an overview of the functional architecture for the SOCIETIES project. The diagram in Fig. 7.3 provides an overview of the platform "core services" provided. These are described at a conceptually very high level to give an understandable overview of the capabilities of the system.

These are further described under a grouping assigned from the logical deployment they support, e.g. services for an administrative domain, for CISs, for CSSs or for every node.

Multi CSS/CIS grouping – these services support a wider open group of stakeholders in a federated domain model. They offer federated search, identity, and domain administration functions, and store multiple CSS or CISs public information. There is one instance of these services per administrative domain, and other federated domains can request information.

This group includes the following services: the *domain authority* (which provides and manages the CSS and CIS identities in a decentralized manner, allowing authentication between multiple domains), the *CIS directory* (which manages the CIS information in a decentralized 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* (which provides search facilities for CSSs, based on their identifier or by specifying search criteria, such as public profile attributes and tags), the *CIS recommendation* (which is responsible for handling CIS recommendations, allowing for

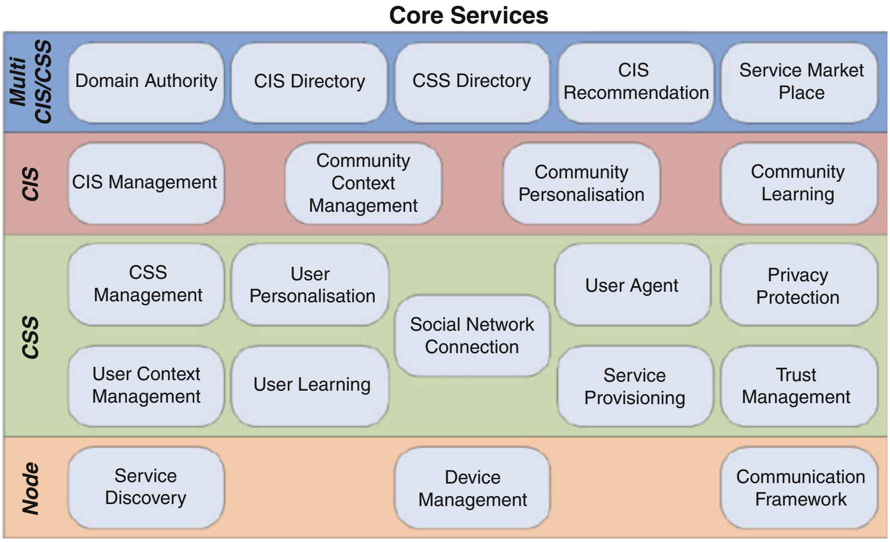

Fig. 7.3 High level architecture of the SOCIETIES platform

recommendations of CISs to users), and the *service market place* (which provides access to a repository of installable third-party (3P) services and optional "core" services, and provides mechanisms for charging for these services).

CIS grouping – these services support a community and their community interaction space (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* [which 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* (which 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* (which supports community preferences and community intent learning), and the *community personalization* (which 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 grouping – these services support a participant and their cooperating smart space (CSS). The word 'participant' is used to *refer to a single user or organization*. 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* (which controls which nodes – devices or cloud instances – are part of the CSS, assigns them a common identity and manages resource sharing and configuration policies)

- The *user context management* (which is responsible for acquiring the user context from sensors and other context sources, for modeling and managing the collected data, for maintaining current and 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 personalization* (which evaluates the user behavioral models, to identify and apply preference changes)
- The *social network connection* [which 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* (which 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* (which supports learning of user behavior models exploiting the user's history of actions stored in the system) and the *user agent* (that acts on behalf of a single CSS based on captured behavior models to establish the system's proactive behavior, and capturing any feedback)
- The *trust management* (which is responsible for collecting, maintaining, and managing all information required for assessing the trust relationships and evaluating direct, indirect, and user perceived trust) and *service provisioning* (which supports the setup and lifecycle control of a 3P service or CSS resource, allowing for installation)

Node grouping – the core services in this group are available per node. A CSS node is a *logical node/device/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* [which provides the necessary mechanisms to support intra- and inter-CSS communication through the discovery of CSS nodes (devices)]
- The *device management* (which provides mechanisms for managing attached devices and management of their capabilities)
- The *service discovery* (which 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)

7.4 SOCIETIES Innovations

In order to realize the SOCIETIES vision and support the three key concepts of discover, connect and organize, for pervasive communities, a number of challenges arise. The primary innovations required to address the challenges will now be detailed.

7.4.1 SNS Interoperability

An important aspect of the pervasive community vision is to capture and facilitate social Web interactions. However we are not producing just another social network. Instead, we need to develop a bridging scheme between the Social Web and pervasive communities.

Interoperability with Social Web initiatives needs to be taken into account, with regard both to standards work on federation and also to non-standardized approaches (e.g., Facebook, Twitter, Google+) given their popularity.

The SOCIETIES Social Connector/Extractor leverages and impacts on Web-based standard APIs and protocols specifically defined for efficient exchange of data among current and upcoming social networking initiatives.

1. *Bridging Social and Pervasive Communities*
   Our aim is to make pervasive communities interact seamlessly with the Social Web (Gallacher et al. 2011). This provides a key bridging or staging point for users not currently using pervasive communities. They can still interact with members of a pervasive community using familiar social networks, while not accessing all the benefits.

   Our first step is to link the profiles and activities of users in our pervasive communities with those in their online social networks (e.g., Facebook, Twitter, Diaspora, etc.). In order to do this, we need a common representation of people, profiles, and activities.

   A key second step is to facilitate exchanges of information between both systems. For example, details such as the name of the pervasive community, the participants' identities, and the activities of individual participants.

Data is exchanged by the following mechanisms:

Pull: SNS information is extracted to inform a participant of what is happening in the Social Web.
Push: pervasive community activity information is updated as user status in the Social Web.
Alignment: automatic and real-time pervasive community profile extraction and updating of corresponding Social Web profiles.

7.4.2 Context-Awareness

In today's systems, no context models have been developed that can support the management of context for a dynamic community in large- scale systems. While research has been carried out regarding community context, there is no support for resolving context conflicts or for inheritance of context information across hierarchical communities. Additionally, while research has been carried out with regard to the formulation of groups of people with common interests, the lifecycle management of these groups does not consider the many ongoing context changes that occur for each user. SOCIETIES is innovating in a number of areas (Doolin et al. 2012) related to this problem space, as described below. These innovations can be related back to the discover, connect, organize paradigm as follows:

Discover: context sensing resources, user context values and available resources based on current context information

Connect: users based on context similarity and connect users with relevant resources based on current and historic context information

Organize: community lifecycles and membership based on context information of individual members and of the entire community

1. *Context modeling*

   The SOCIETIES community context model (CCM) assists in the management of dynamically formed communities of people, crossing the digital and physical worlds, in which users discover relevant entities. Context information is one of the main criteria that is used to identify which users are relevant to which individuals, resources, and communities.

   SOCIETIES CCM allows modeling of context-related information for: communities of individuals, quality of community context, history of community context, social relationships (and social media information in general), bonds and relationships between users, and interactions among users.

   New context modeling approaches (Roussaki et al. 2012) have been developed in order to properly represent community context information. In this respect, hybrid models are used that support a multitude of community context representations, based on the type of the respective information, its quality and source. Thus, both deterministic and stochastic solutions are adopted, which vary depending on the qualitative and quantitative inherent features each context type may demonstrate. Furthermore, the conflicting context information that may be observed within communities is addressed with a flexible model that is capable of supporting context inheritance and refinement mechanisms.

2. *Community context extraction*

   Estimation of community context must be carried out both on demand and continuously. SOCIETIES context estimation can provide useful real-time information for community activities, and can connect people with common interests or other context commonalities/bonds. This forms a key part of the "relevance" aspect of the SOCIETIES discover, connect, organize paradigm. There are a number of models that are currently used to represent and estimate these community context values, such as: aggregation of the values of user context data for community members; stochastic representation of this mainly for discrete or enumerated context value formats, average/median values for the discrete or enumerated context value formats, most probable value, etc. The community context extraction is a difficult task that becomes more challenging as the complexity and dynamicity of context information rises and as the number of the community members increases.

3. *Sources of context updates*

   In order to update context information, SOCIETIES specifies a number of context information sources that can capture context values. Individual context data updates can be triggered by sensors, user actions, changes in social media data, quality of context thresholds, context refinement (inference) processes, community activity, community membership changes, etc.

As an example, a change in an individual user's context information may trigger the update of one or more pieces of community context information across multiple communities of which that person is a member. This process of keeping a community's context constantly updated enhances its usefulness, relevance, and trustworthiness, thus allowing more people to connect to it and hence to each other.

4. *Context similarity*
Context similarity evaluation (CSE) in SOCIETIES is a challenging research area that requires new algorithms/mechanisms for context comparison. It is necessary to enable the evaluation of context similarity for both quantifiable context information such as location coordinates, weight, or temperature (which require arithmetic methods), as well as for qualitative context information such as user interests, status, or symbolic location (which require new principles in comparing context semantics to estimate similarity).

The CSE concept supports a number of decision-making components within the SOCIETIES framework for tasks, such as: creation/discover/deletion of communities, identification of community hierarchies, identification of communities that could be merged/split, community membership management (i.e., addition/removal of members), prediction of user intent based on similar context (as common context often indicates similar intentions), and user preference discovery based on similar context (as the same preferences may be applicable not only under a given combination of context information values, but also when similar context conditions are observed).

5. *Location inference and prediction*
Location is a core part of context information, and the techniques used to infer and predict location in SOCIETIES deserve some attention. Current approaches usually require manual check-in (e.g., Facebook, FourSquare) or semi-automatic check-in (e.g., Google Maps/Latitude). SOCIETIES focuses on automatic estimation of indoor location (through the use of Wi-Fi sniffers for example) with fine granularity (as opposed to the common practice of the manual check-in seen in many of today's mobile services). In distinction to other competitive solutions that support automatic location tracking, our approach does not require any expensive hardware. Furthermore, SOCIETIES provides an outdoor location prediction framework that exploits cellular traces and communication records, and forecasts a user's locations in a given time frame based on periodicity.

7.4.3 Learning

Learning in a social media context tends to be mainly for the benefit of the service provider, for example data mining for revenue generation through targeted advertising, whereas in a pervasive system, learning focuses more on the user by implicitly acquiring information which would otherwise have to be acquired explicitly – and typically only individual users are considered by such systems. SOCIETIES surpasses this current impasse by learning on a community level to underpin many of the project innovations at various points in the platform.

In previous research projects [such as PERSIST (2012) or DAIDALOS (2012)] preference learning mechanisms focused on the individual user. In SOCIETIES, we break new ground by learning about preferences of communities of users. This learning is based on historic behavior and context information collected from all members of a community (bearing in mind the privacy requirements of users) and fused together to create a single behavior and context history for the entire community. Context-dependent community preferences are extracted with these learning techniques, and are associated with a related community. Current and new community members can inherit all or part of the learned/modified community preferences to enhance their own preference set.

In terms of our DCO concept, learning relates as follows:

Discover: individual and community preferences
Connect: learning supports connection based on individual/community preferences
Organize: community-level learning assists individuals in acquiring information and links from other community members.

7.4.4 Privacy

Within SOCIETIES, all design has privacy built in as a fundamental requirement. Our "privacy-by-design" approach means that privacy protection is intimately integrated in the platform, rather than being an afterthought. SOCIETIES privacy protection supports the user in management of personal data (Gonçalves et al. 2012), enforcing negotiations between user preferences and service/community policies, data disclosure with personalized obfuscation, multi-identity management and selection, and privacy assessment.

In relation to DCO, the following can be stated:

Discover: individuals and communities who will comply with a user's privacy preferences.
Connect: privacy policy negotiation during connection with potential for data obfuscation, micro-agreement, and a privacy-aware social firewall
Organize: privacy audit/assessment contributes to reputation, and this enables more informed organizational activities.

1. *Personalised privacy policy negotiation*
   Privacy policy negotiation provides users with the ability to choose the personal data they wish to disclose to other entities. SOCIETIES uses user preferences to automate the privacy policy negotiation process, which in turn benefits the user by lifting the burden of privacy policy configuration from them. This also ensures that the data owner and requestor are bound by a common data disclosure

agreement. Policies also help the user to stay consistent in terms of the data they disclose, as the same preference can map to numerous situations. Privacy policies can be used to allow the user to assess suggested communities (automatically or manually), and privacy policy negotiation within SOCIETIES communities allows tailored data disclosure based on user needs.

2. *Data obfuscation*
   The purpose of data obfuscation is to reduce the personal content of data. The concept is an amalgamation of two "privacy-by-design" principles – data minimization and enforcement. While location and identity have been obfuscated in other research projects, SOCIETIES formalizes this obfuscation, generalizes it and applies it to many data types. When sharing data on behalf of a user, obfuscation takes place before any data are shared with others.
3. *Personalization and learning for privacy*
   SOCIETIES personalization and learning techniques are applied to data disclosure to support the learning of privacy preferences, and to facilitate automation of negotiation processes, identity selection, and obfuscating personal data.
4. *Micro-agreements for business*
   In support of enterprise, SOCIETIES support for micro-agreements allows employees to expose their company privacy policies for the social and pervasive mobile services they wish to consume or offer using micro-policies. Automatic negotiation of micro-agreements makes the connection between new community service providers and the service consumers in a secure and trustworthy manner. Companies can use the content of micro-agreements to monitor quality of service provided by service providers, and can organize their collaboration with other parties accordingly.
5. *Privacy assessment: privacy-aware social firewall*
   This novel area of research within SOCIETIES protects users from unintentional privacy leakages into social channels by monitoring data flows, estimating privacy risks, and passing control of them over to the user. The firewall operates based on privacy assessment which includes the following elements: auditing actual privacy practices by monitoring and logging how private information is actually used by applications and services, detecting privacy breaches and/or detection of potential privacy information leakage, and enforcing privacy by requesting immediate user feedback in case of potential violations. This usage of privacy assessment generates trust-relationship information, and is consequently an enabler of trustworthy entity discovery and connection.

7.4.5 Trust

Interaction between users, communities, and services requires formal trust assessment including robust and authenticated mechanisms. Aspects of trust (e.g., purposeful or referral) and trust relationships are the subject of current research, but these are not used in social media. Trust-related research supports the SOCIETIES DCO paradigm as follows:

Discover: trustworthiness of individuals, communities, services, and entities that you trust in advance.
Connect: to various entities based on individual and community trust assessment.
Organize: trust-based community membership management and trust-based community lifecycle (merging and splitting) based on trust relationships among members of existing communities.

Community-enhanced trust assessment based on feedback and learning
Assessment of the trustworthiness of individuals, communities, and services in SOCIETIES is based not only on the experiences of the user, but also on the experiences of fellow members of a community. Inherent in this process are SOCIETIES' feedback and learning mechanisms. This assessment provides the necessary facilities for trust-based discovery of individuals, communities and services. It supports automatic interconnection of, and information sharing of privacy-sensitive information with, trusted entities. SOCIETIES also supports community lifecycle and hierarchy management, based on the trust relationships among members of the parent community, and trust-based community lifecycle transition from temporary to ongoing, and facilitates the trust assessment process with community feedback. The SOCIETIES trust management and evaluation framework provides the necessary support infrastructure for the maintenance and management of dynamically changing trustworthiness within collaborative domains.

SOCIETIES researches direct trust (which is evaluated based on the interaction history of users and trust ageing) and indirect trust (in which case trust is inferred based on a user's fellow community members). A fusion of these trust types allows assessment of the aggregate trust value as perceived by the user. More specifically, some people employ objective measures to evaluate their level of trust in an entity, while others rely on a more subjective feeling. Direct trust generally outweighs indirect trust in this fusion process. However, the weight of each factor also depends on the confidence level with which it has been estimated.

7.4.6 Community Orchestration

Existing social networks provide a means to organize social network connections into communities (groups, circles, or cliques). However, the management of communities in current social media is largely manual, with some assistance provided in the form of suggestions.

In addition to the "connection" management overhead of the social interactions, Pervasive communities present a completely new challenge due to the physical element (devices, sensors, locations, and services) captured in the pervasive community concept. This adds "real world" impact to a group or community.

In order to maximize the value for users, we need to allow users to organize all aspects of communities. *Orchestration* refers to the ability to manage the intelligent formation, organization, membership, and termination of communities. Relating this back to the value propositions of the SOCIETIES project:

Discover: potential communities and members.
Connect: individuals via community membership/formation.
Organize: individuals into communities, form, merge, and delete communities and sub-communities.

1. *Context state model*
   Individuals and communities generate large amounts of data continuously, and this data needs to be gathered, with due regard to individuals' privacy policies, and processed so that it can be analyzed in a time-ordered manner. Potential new communities, defunct communities, and existing communities which an individual might wish to join/leave are identified via context state models (CSM).

   Each context state model (CSM) is made up of key attributes that describe characteristics relevant to existing and potential communities. Once a CSM has been created or modified, locality-sensitive hashing (LSH) is used to reduce the CSM to a simple rapidly computable value. The CSM data of individuals and communities is compared to that of others and trajectory mining techniques are used to discover similarities.

   CSMs can be used for analysis in near-real-time group dynamics. CSM modeling allows quick computational comparisons, which is vital for the discovery process, to ensure that a user has a 'live' (near-real-time) experience. CSMs also permit user data to be analyzed anonymously in an abstract fashion. A *community life-cycle manager* enforces the individual privacy constraints of users prior to any automatic or semi-automatic orchestration activity. The process is event-driven, and reacts to events as they occur to determine how group dynamics are affected by CSM state changes.
2. *Community nature*
   Community nature refers to the concept of "temporary" and "ongoing" communities. Some communities form and can last for years, while others are formed on demand and are discarded within minutes. For example, a community formed around a family will have significant longevity, but one formed among a group of people waiting at a bus stop is likely to be short-lived. The ability to distinguish between communities that are likely to be short-lived (e.g., location- or purpose-based) and those that are likely to be of indefinite duration (e.g., family-based) can assist in efficient community orchestration and life-cycle management.

   Algorithms that are specific to the ongoing/temporary nature of communities can autonomously/semi-autonomously drive the creation, configuration, and deletion of communities on behalf of the user. This relieves the user of lots of housekeeping tasks associated with the many communities to which they belong during their daily lives.

The community nature, i.e., the concepts of "ongoing" and "temporary", can be exploited by the system to assist an end user. Longevity of association between individuals is clearly a key driver for potential community formation and, in the online world, it may be acceptable for this to be a prerequisite. However, in the physical pervasive world, we need also to support ad hoc community formation for short-term goals which would not be possible with a longevity prerequisite. Discovery algorithms can look for brief, temporary communities at short intervals and longer-lasting and ongoing ones at longer intervals. The data used for this and how to analyze it varies depending on how long a community might be expected to last. Temporal and community nature semantics are attached to communities, and used to inform decisions such as deleting obsolete communities, configuring them, and creating new ones.

7.4.7 User Intent

The actions of people are invariably guided by a set of tasks that they need to complete over the short to medium term, e.g., buying fresh milk from the local shop. Ideally, for a system to fit in with the user's view of the world, it also needs to be aware of what the user is trying to achieve, i.e., their final goal. We refer to this as user intent.

User intent is not formally considered in social media. Instead they focus on the lower level of social interactions, without considering why these interactions take place. Some systems rely on manual entry of the tasks and subtasks, and require the user to continuously update these tasks with progress. In SOCIETIES, we attempt to capture intent by observing the user, and attempting to deduce what their intentions are.

Most user intent learning approaches aim to predict future user intentions based on raw context data. This is very difficult due to the gap between the low semantic level of raw context and the high semantic level of intentions.

User goals have been considered in pervasive systems, but predicting future behaviors has been limited to the level of the individual and not to a community. Relating this to the value proposition of SOCIETIES:

Discover: proactively discovers both individual and community intent.

Connect: once the intent of different users (e.g., a community) is discovered, we are able to effectively connect them by establishing a community with users who share similar intent.

Organize: provides community intent-aware services, and takes actions on behalf of a community.

1. *User intent prediction based on CRFs*
   Rather than directly using raw context data to learn intents, we introduce the concept of a situation to fill in the semantic gap. A recurring pattern in the context information yields a specific "situation" associated with a clique of raw context data.

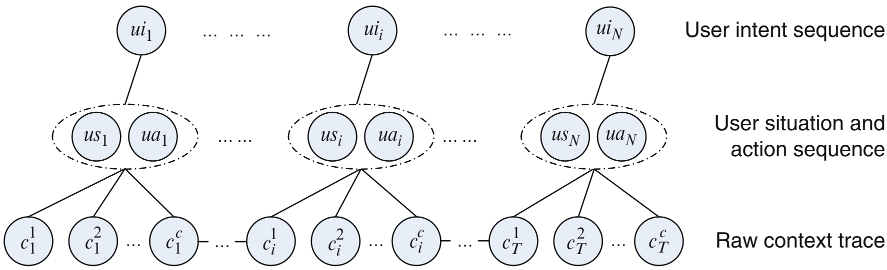

Fig. 7.4 Prediction based on CRFs

Then, based on the situation, as well as the real-time user action, user intent discovery and prediction are performed. Specifically, a 3-layer conditional random field (CRF) is constructed to model the relationship among the raw context data, the situations/actions, and the users' intentions, as shown in Fig. 7.4.

This model is capable of facilitating social interaction by proactively learning/detecting/predicting both individual and community intentions.

2. *Context-aware user intent prediction*

User intent is also predicted in SOCIETIES via a second approach, which also exploits context information but employs methods of statistical analysis. In this approach, the user behavior is modeled based on the user interactions with services. Text compression algorithms are used for identifying sequences of *user actions* and *user tasks*. Each user action is accompanied by a context data snapshot describing the situation of the user. Frequent occurring sequences of user actions, and their associated context snapshots, are grouped into user tasks. The user behavior model is described by stochastic models (variable-order Markov chains) that describe the sequences of user tasks and user actions and the respective transition probabilities.

History logs containing individual user interactions, along with context information are collected in a common context history repository. Alternatively, individual intent models can be provided. Learning algorithms, also used for the discovery of the user intent model of individuals, create a community-wide behavior model. The community model is used in order to enhance the intent prediction for individual users. Figure 7.5 provides an example of this concept.

Case A outlines a scenario where the history logs containing individual user's interactions, along with corresponding context information, are collected in a common context history repository. This common context history repository is used to discover a community intent model.

Case B outlines a scenario where individual intent models, which have been prior-discovered from the individuals' history logs, are provided as a substitute for the common context information, i.e., the community intent model is discovered from the aggregation of the individual intent models.

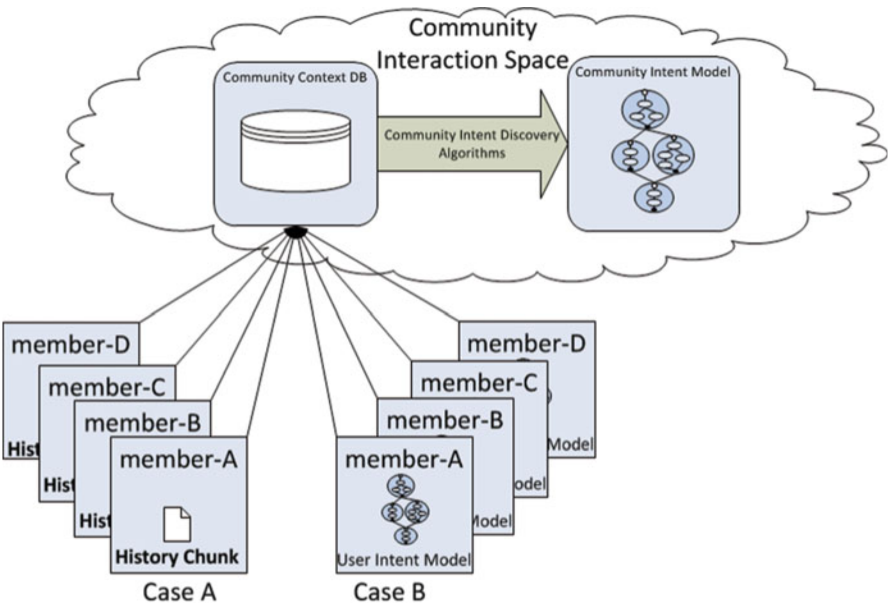

Fig. 7.5 Community intent model learning process

7.5 SOCIETIES Methodology

The SOCIETIES platform is intended to be used by individuals and communities. It is therefore imperative that the design is user-centered and, wherever possible, user-driven. The SOCIETIES methodology (Jennings et al. 2011) is based on the inclusion of three different user communities, who have engaged continually in the design and development process from initial user requirements gathering and early concept creation through to scenario refinements via early prototype evaluations. These are an *enterprise* (Lima et al. 2012) or *business* community, a *student* (Gallacher et al. 2012) or *young technology-wise* community, and a *disaster management* (Floch et al. 2012) or *emergency rapid response* community.

Early user requirements were elicited via ethnographic techniques, involving close observation of each community in their real-world environment, combined with questionnaires and participatory discussions. Requirements and scenarios were further refined through users' engagement with storyboards, and in the case of the students, a Wizard-of-Oz experiment. In the next steps of the project's user evaluation, each user community will evaluate the advanced system prototypes in situ in their real-world environments.

The *student community* is comprised of computer science students, for the most part avid social media users and naturally comfortable with many digital devices. They envision many potential uses for pervasive community systems, such as enhancing management of access to common services, such as libraries, and supporting collaboration both among students themselves, as well as between students and their lecturers. Students collaborated with developers through participatory workshops; using brainstorming, bodystorming, and storyboards to illustrate and investigate scenarios, while some students also opted to participate in a Wizard-of-Oz experiment that examined user acceptance of some proposed features and services of pervasive communities set in a campus environment. The results from the early student trials indicate that in general the students see value in joining pervasive communities. Clear controls over privacy settings and automation of services are a prerequisite for use of such a system according to this student community, but once these are in place, they are open to sharing some personal data and preferences, in order to add value to and leverage value from those communities.

Enterprise users are primarily interested in pervasive community services, which could enhance the workflow of their daily lives, with additional features for more selective relevant communications. These busy executives, in several cases, report being overloaded with communications and information, and the automatic filtering of both, without missing opportunities for useful connections, is attractive to them. Online community group discussions, observation, and storyboard-based workshops have supported the project's engagement with this group, and have led the project to focus on pervasive communities in a conference scenario. Benefits for enterprise attendees could potentially then include services such as personalized real-time information, such as agendas and navigation aides, events tailored to registered attendees' interests, live feedback during presentations, and maximizing networking opportunities. Many enterprise group executives indicate willingness to share some of their preferences with the system, while privacy, with regard to which preferences are shared for services, is clearly also a requirement.

The experienced *disaster management* community, with which we are engaged, is interested in any means of communication that supports the relief effort in a disaster zone. While emergency disaster situations are unpredictable by nature and difficult to envision, the project engaged with this group, composed of professional individuals from several different countries, during intensive assessment mission field training exercises, using mixed methods observation, questionnaires, discussions, and storyboard workshops. Relevant detailed and reliable information is important to disaster missions, and these disaster management professionals recognize great potential value in SOCIETIES providing this via pervasive communities and crowd computing. They are interested in tapping the expertise and knowledge of wider society in a disaster situation, for particular tasks, once that information can be filtered and effectively validated. They see advantages in the system coordinating the broad population who would like to help to assist a disaster mission, for certain tasks, such as checking satellite imagery for infrastructure damage, translation, and transcribing logbooks. Our research findings show it is important for disaster management users to feel in control of the system, and that they would prefer the system to make intelligent suggestions to support humans making decisions, rather than automation for most tasks.

In summary, the 'voice of the user' has been considered and continually integrated into the system design, through user-centered design and participatory methods. The next steps in the project will be to immerse the three user communities into prototypes of pervasive environments that have been derived and built from the scenarios and storyboards presented to them in the previous trials, and to evaluate their responses to these pervasive environments.

7.6 SOCIETIES Status, Future, and Exploitation

At the time of writing, the SOCIETIES project has completed a number of phases, and is now focused on initial implementation and integration of the platform described above.

Requirements and scenarios have been developed and verified with our user groups (SOCIETIES 2011), which encapsulate the critical use-cases of the project. Requirements have been specified in three specific domains – user requirements, technical requirements, and business requirements – each of which is combined to produce the overall requirement specification of this complex system. The requirements and scenarios for the project have been submitted to an initial user evaluation for verification and feedback (SOCIETIES 2011).

The SOCIETIES' system architecture has been specified along a number of interrelated lines of research: the overall system architecture that defines the high level architectural features of the SOCIETIES system, the service architecture that examines infrastructural and end-user service requirements, and finally interoperability architecture that specifies how SOCIETIES will integrate beyond its research remit.

The project core platform and intelligent enabling services have been designed in detail based on the architecture and requirements specified in the early phases of the project, and cover all aspects of the system design from client frameworks, communications management, device management, and service lifecycle management, to intelligent community orchestration, personalization, context management, and privacy & trust.

Development and integration of the SOCIETIES platform is driven by our commitment to carrying out trials with real users by the end of 2012, followed by a further research, development, and trial phase in 2013.

Looking to the future, the SOCIETIES project outputs (SOCIETIES 2011) can be exploited in many ways, from development of a full commercial community management system to integration of various enabling technologies into existing and emerging products, and to development of an open-source community around the results. SOCIETIES' stakeholders, identified in our business models, range from the mass-market to service providers, and onwards towards a number of secondary customer groups such as conference centers, emergency organizations, and universities.

In terms of mass-market adoption, SOCIETIES has involved real user communities throughout all stages of our research, from initial concepts and scenarios to fully developed software. The purpose of this is to identify real consumers' needs and wants, and their level of preparedness to adopt this highly pervasive technology.

Other stakeholders include brokers, who can harness the rich data produced by the SOCIETIES platform and offer it to various service providers. Service consumers, i.e., end users, require solutions to their discovery needs, their device interoperability and connection requirements, and the ability to bridge the gap between the real and virtual worlds. Service providers will have the capability to create new highly intelligent community-based services, as well as enhancing existing offerings with rich relevant information. Social network service providers can increase their SNS capabilities through the integration of SOCIETIES' enhanced discovery results, through the integration of connected communities, and through the provision of context-aware, relevant physical-world data to end users.

As noted previously, privacy is a major concern for many consumers. Consumers need to know that their data will be protected, and which communities have relevant trust levels based on personal requirements.

To support the adoption of SOCIETIES' results, the project is committed to providing external interfaces (APIs) to allow third-party entities to exploit all or parts of the system. The establishment of trust and other required customer relationships will facilitate the creation of a strong user base, as will integration with existing technologies. SOCIETIES is promoting itself to the Future Internet community, through various dissemination channels, in order to stimulate debate and discussion on the results. To increase consumer confidence in the SOCIETIES system, demonstrating a capability to monitor (in a privacy supporting manner) and learn about users in order to take proactive actions on their behalf is a key requirement. Proactive adaptation of a community's characteristics based on context, learned information, individual and group preferences, and membership management will enhance users' confidence in the value of the project results. Finally, the provision of exemplar community-based services, which do not exist in today's market, will further stimulate adoption.

7.7 Summary and Conclusion

In summary, SOCIETIES intends to bridge the gap between new social media and the virtual and real worlds, driving the adoption of pervasive technology through the discovery, connection, and organization of purpose-driven communities of interest. SOCIETIES believes that the individual elements of *discover*, *connect*, and *organize* make for a compelling overall value proposition that provides identified stakeholders with real value in today's competitive marketplace. SOCIETIES merges concepts from pervasive computing and social computing, breaking the physical–digital boundary, in a way that ensures the privacy of the user is safeguarded, letting users control their data, which in turn builds user trust.

Acknowledgments This work was supported in part by the European Commission as part of the SOCIETIES project under the FP7 programme which the authors gratefully acknowledge. The authors also wish to thank our colleagues in the project, without whom this paper would not have been possible (SOCIETIES 2011). Note that apart from funding this project, the European Commission has no responsibility for the content of this chapter.

References

DAIDALOS, (2012). *Project Website*. http://www.ist-daidalos.org. Accessed 16 Nov 2012.

Doolin, K., Roussaki, I., Roddy, M., Kalatzis, N., Papadopoulou, E., Taylor, N. K., Liampotis, N., McKitterick, D., Jennings, E., & Kosmides, P. (2012). SOCIETIES: Where pervasive meets social. *Future Internet Assembly Book, 1*, 30–41.

Floch, J., Angermann, M., Jennings, E., & Roddy, M. (2012). *Exploring cooperating smart spaces for efficient collaboration in disaster management*. Vancouver: ISCRAM.

Gallacher, S., Papadopoulou, E., Taylor, N.K., Blackmun, F.R., Williams, M.H., Roussaki, I., Kalatzis, N., Liampotis, N., & Zhang, D. (2011). Personalisation in a system combining pervasiveness and social networking. In *20th International Conference on Computer Communications and Networks (ICCCN 2011)*. Maui.

Gallacher, S., Papadopoulou, E., Taylor, N.K., Blackmun, F.R., & Williams, M.H. (2012). Intelligent systems that combine pervasive computing and social networking. In *9th IEEE International Conference on Ubiquitous Intelligence and Computing (UIC 2012)*. Fukuoka.

Gonçalves, J., Gomes, D., & Aguiar, R. (2012). Low-latency privacy-enabled context distribution architecture. In *IEEE Communication Software Services and Multimedia Applications Symposium (IEEE ICC 2012)*. Ottawa.

Greenfield, A. (2006). *Everyware: The dawning age of ubiquitous computing*. Berkeley: New Riders.

Hansmann, U., Merk, L., Nicklous, M., & Stober, T. (2012). *Pervasive computing: The mobile world* (2nd ed.). New York: Springer.

Jennings, E., Madden, E., Roddy, M., & Roussaki, I. (2011). Participative user research and evaluation methodologies for pervasive communities. In *Irish Human Computer Interaction Conference 2011 (iHCI 2011), Integrated Practice Inclusive Design (CIT)*. Cork.

Lima, C., Gomes, D., & Aguiar, R. (2012). Pervasive CSCW for smart spaces communities. In *IEEE International Conference on Pervasive Computing and Communications Workshops 2012 (PERCOM Workshops 2012)*. Lugano.

Lukowicz, P., Pentland, S., & Ferscha, A. (2012). From context awareness to socially aware computing. *IEEE Pervasive Computing, 11*(1), 32–41.

Obaidat, M. S., Denko, M., & Woungang, I. (2011). *Pervasive computing and networking* (1st ed.). Hoboken: Wiley.

Pallis, G., Zeinalipour-Yazti, D., & Dikaiakos, M. D. (2011). New directions in web data management. *Studies in Computational Intelligence, 331*, 213–234. doi:10.1007/978-3-642-17551-0_8. Online Social Networks: Status and Trends.

Parthasarathy, S., Ruan, Y., & Satuluri, V. (2011). Community discovery in social networks: Applications, methods and emerging trends. In *Social network data analytics* (pp. 79–113). Springer US. doi: 10.1007/978-1-4419-8462-3_4.

PERSIST. (2012). *Project Website*. http://www.ict-persist.eu. Accessed 16 Nov 2012.

Roussaki, I., Kalatzis, N., Liampotis, N., Kosmides, P., Anagnostou, M., Doolin, K., Jennings, E., Bouloudis, Y., & Xynogalas, S. (2012). Context-awareness in wireless and mobile computing revisited to embrace social networking. *IEEE Communications Magazine, 50*(6), 74–81.

Semertzidis, T., & Laso-Ballesteros, I.D.P. (2010). *Social networks: Current trends and research challenges, coordinated by the next MEDIA coordination action*.

SOCIETIES Deliverable D2.1 *Specification of initial scenarios and user requirements*. http://www.ict-societies.eu/files/2011/11/SOCIETIES_D21.pdf. Accessed 16 Nov 2012.

SOCIETIES Deliverable D2.2 *Scenario description, use cases and technical requirements specification*. http://www.ict-societies.eu/files/2011/11/SOCIETIES_D22.pdf. Accessed 16 Nov 2012.

SOCIETIES Deliverable D8.1 *Paper trial evaluation report*. http://www.ict-societies.eu/files/2011/11/D8.1_public.pdf. Accessed 16 Nov 2012.

SOCIETIES deliverables. http://www.ict-societies.eu/project-deliverables/. Accessed 16 Nov 2012.

SOCIETIES Integrated Project. *Full title: Self Orchestrating CommunIty ambiEnT IntelligEnce Spaces*, http://www.ict-societies.eu. Accessed 16 Nov 2012.

SOCIETIES project partners. http://www.ict-societies.eu/partners/. Accessed 16 Nov 2012.

Xu, B., Chin, A., Wang, H., Chang, L., Zhang, K., Yin, F., & Zhang, L. (2011). Physical proximity and online user behaviour in an indoor mobile social networking application. In *Internet of Things (iThings/CPSCom), 2011 International Conference on and 4th International Conference on Cyber, Physical and Social Computing* (pp. 273–282), doi: 10.1109/iThings/CPSCom.2011.74.

Zaphiris, P., & Ang, C. S. (2009). *Social computing and virtual communities*. Hoboken: Chapman and Hall/CRC/IGI Global.