Pervasive and Mobile Computing 9 (2013) 437–453

Contents lists available at SciVerse ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Review

Mobile social networking middleware: A survey

Paolo Bellavista[a,*], Rebecca Montanari[a], Sajal K. Das[b,1]

[a] University of Bologna, Viale Risorgimento, 2 – 40136 Bologna, Italy
[b] The University of Texas at Arlington, P.O. Box 19015, Arlington, TX 76019, USA

ARTICLE INFO

Article history:
Received 26 July 2011
Received in revised form 4 March 2013
Accepted 9 March 2013
Available online 21 March 2013

Keywords:
Mobile social networking
Middleware
Social communities
Mobile applications
Opportunistic computing

ABSTRACT

The convergence of social networking and mobile computing is expected to generate a new class of applications, called *Mobile Social Networking* (MSN) applications, that will be of significant importance for the coming years. Indeed, MSN enhances the capabilities of more traditional Online Social Networking (OSN) to a great extent by enabling mobile users to benefit from opportunistically created social communities; these communities should be determined not only by common interests or contacts but also by mobility-related context, such as physical location and co-presence. In this paper we precisely define what we intend for MSN applications and overview the primary MSN support solutions available in the current literature and that specifically address the underlying technical challenges, design issues, and emerging middleware guidelines. Our primary goal is to identify engineering design criteria for future MSN middleware solutions, capable to flexibly adapt to different application domains and deployment requirements. To this purpose, we present a novel taxonomy of MSN structures and describe how various existing middleware approaches fit the proposed classification; moreover, the survey takes the opportunity of these descriptions to discuss related middleware design/implementation choices determining specific tradeoffs between expressive power, flexibility, and scalability.

Contents

* Corresponding author. Tel.: +39 051 2093001; fax: +39 051 2093073.
E-mail addresses: paolo.bellavista@unibo.it (P. Bellavista), rebecca.montanari@unibo.it (R. Montanari), das@cse.uta.edu (S.K. Das).
1 Tel.: +1 817 272 7405; fax: +1 817 272 3784.

1574-1192/$ – see front matter
http://dx.doi.org/10.1016/j.pmcj.2013.03.001

1. Introduction

The success and widespread diffusion of Online Social Networking (OSN) applications in our daily lives are being manifested in various ways. At the same time, the exploding penetration and capabilities of smart phones have already made the mobile service market a priority for the industry. For instance, let us think about how mobile applications are pushing further the production of user-generated contents anywhere and anytime. Similarly, consider how traditional OSN is stimulating the sharing of an increasing amount of personal information and resources, even by changing the common perception of privacy-related boundaries. It is somewhat unavoidable that the two worlds of mobile applications and social networking will gain increasing importance and their widespread diffusion will lead to important synergy and convergence. Example scenarios abound, in both the scientific literature and the mass market, from mobile Web access to traditional OSN applications, to more interesting and novel application scenarios related to crowdsourcing-based participatory sensing via mobile phones [1] as well as opportunistic collaborative networking/computing that takes advantage of social awareness [2,3].

In the following we clearly distinguish between traditional OSN solutions/applications and the recent concept of Mobile Social Networking (MSN) [3], which relates to the novel relevant opportunities opened by the full exploitation of mobile and opportunistic computing techniques, e.g., to dynamically determine temporary social communities, also based on local considerations, physical location, and co-presence, as better detailed in the following. Let us note that we consider the mobile access (via smart phones and their Internet connectivity) to traditional social networking applications, such as Facebook, Twitter, or LinkedIn as they are nowadays, as a case of OSN systems/applications; they are less interesting from the novelty point of view and out of the scope of this paper. In contrast, we focus on MSN solutions that enable the building of very dynamic and context-aware social groups, possibly by maintaining/processing social metadata also locally, without the need of continuous connectivity with a fixed Internet server infrastructure.

Some interesting approaches, especially from academic research, about mobile opportunistic networking/computing that exploits social awareness as an additional source of context metadata, have recently been published [1–6]. However, most of them are not concentrated on offering specific abstractions for mobility-oriented social networking and/or on the middleware support of novel MSN applications; they are generally more oriented to exploit available social metadata (typically from traditional OSN) with the aim of optimizing mobile communications, networking, and resource management. For instance, the recent literature includes several interesting and impactful solutions in the areas of social-aware collaborative sensing and content distribution/routing [1,7]. In contrast, our goal in this paper is to offer an overview of the major middleware-oriented research trends for the support of novel forms of MSN applications, fully benefitting from the opportunities raised by mobile and opportunistic computing. For instance, we aim to address important yet open questions, such as how does the convergence of mobile applications and social networking enable new forms of social interactions based on physical proximity, co-presence, and frequency of encounters? Or how is it possible to design and implement effective middleware supports to enable novel MSN applications that go well beyond the traditional client–server interaction (where a mobile node accesses social metadata from a centralized OSN server)?

A first survey specifically addressing MSN applications, network architectures, and protocols, appeared recently in [8], thus demonstrating the high level of interest in the field. Differently from that paper, our article does not provide readers with a general overview of algorithms and protocols proposed to solve MSN-specific issues (e.g., from community detection to content distribution, from mobility modeling to security management). On the contrary, we propose a novel taxonomy of possible architectures and organizations for MSN systems and middleware supports, by clearly pointing out their underlying characteristics, merits, and demerits. This taxonomy allows us to derive a series of engineering design criteria and guidelines for application/deployment-specific MSN middleware, capable of achieving the most suitable tradeoff between management overhead, completeness of social awareness, and ability to promptly/locally adapt. Finally, the paper tries to clearly identify the primary MSN open issues and related research opportunities, thus delineating a possible roadmap of priorities to be addressed in the near future.

2. The MSN paradigm shift and a classification of MSN structures

It is widely recognized that a social structure usually characterizes individual's life: people go to bars and restaurants, study together in schools, and work in teams on production lines and in almost every business. Every person plays a role within a social community. Since the 1960s, research and industry efforts have investigated advanced collaborative systems for leveraging human connections and improving human interactions in workplace environments. Although the field of computer-supported cooperative work has contributed significantly to this direction, these solutions generally focus on interactions driven by business needs, where connections among people tend to be formal and structured. Only recently has the convergence of social and computing disciplines focused attention to the design of OSN applications. Their usefulness has been widely perceived, because of their ability to support human social interactions by freeing interpersonal interactions from geographic constraints, thereby enabling anywhere and anytime communication. OSN applications, such as Facebook, MySpace, and LinkedIn, enable online social interactions and lead to a *paradigm shift from physical to virtual communities*, by promoting a Web of social bindings independently of the current physical location of social network members. In other words, social networking applications have played the role of powerful enablers of a social shift, from physical to virtual communities, in the way all of us communicate and interwork in today's world. Part of this fascination closely relates to the

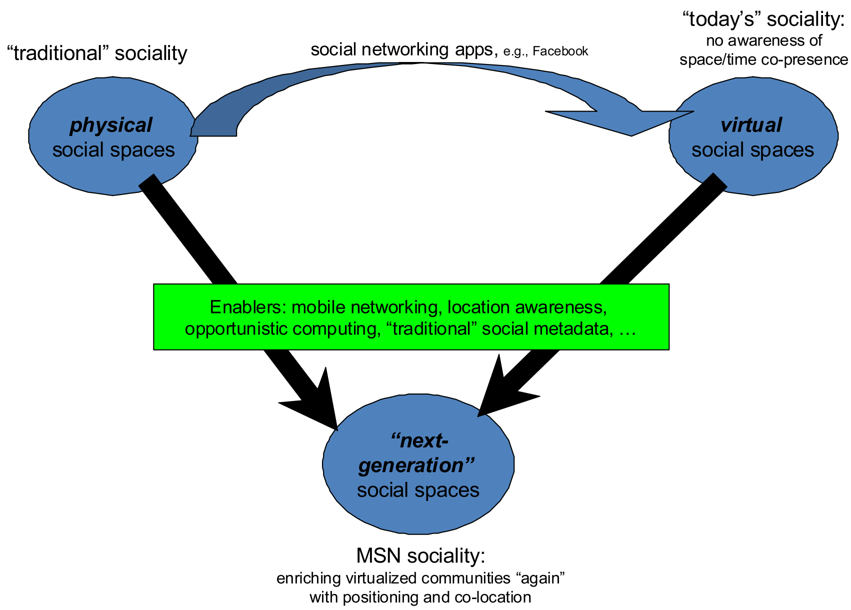

Fig. 1. From physical to virtual to "next-generation" social spaces.

possibility of interacting without constraints for users to be co-located in time and/or in space, thus opening a new, ever growing world of opportunities for our social interactions.

But at the same time the emerging trend toward cyber-physical convergence is opening tremendous new opportunities for pervasive computing, especially in terms of social-aware opportunistic networking and applications, from collaborative opportunistic sensing to community-oriented resource sharing and pervasive healthcare [3]. In particular, in this paper we discuss the emergence and characteristics of a new generation of social applications enabled by today's mobile and opportunistic computing, denoted in the literature as MSN applications to clearly distinguish them from more traditional OSN. As a key feature, MSN applications exploit full awareness of physical positioning and co-location, together with partial storage and ownership of social metadata on users' nodes, to enable new forms of dynamic and physically-bounded virtual communities [9,10]. For instance, in MSN applications a user can bring with her a partial/total knowledge of her social relationships, together with other context metadata such as preference profile and current location; she can exchange (part of) this information with other users depending on mutual positioning and simultaneous proximity to recognized places of social interest, in order to detect and build a temporary and dynamically personalized virtual community, e.g., interested in exchanging messages, opinions, and recommendations. On the one hand, these "next-generation" MSN virtual communities can be realized via (temporarily) sharing context metadata on a server infrastructure that builds personalized views of the social communities of interest. On the other hand, the same goal is achievable in a completely decentralized way by processing and maintaining these views locally to the users' personal devices, with interesting advantages also in terms of fully distributed ownership of social metadata. It is recognized that the possibility of creating ad-hoc social networks between proximate users with common interests promotes anywhere anytime serendipitous social encounters characterized by swarming, transitory, and informal qualities [10,11]. Let us point out that the further paradigm shift we are experiencing is from totally virtualized social communities (e.g., most virtual spaces implemented by OSN applications today are agnostic to positioning and co-location), to location/proximity-dependent virtual communities, capable of exploiting cyber-physical convergence, as illustrated in Fig. 1.

2.1. A classification of MSN structures

MSN applications exhibit unique characteristics that raise several management issues, such as community detection and inference of new social relationships (see Section 3). Although substantial research exists about the design, implementation, and evaluation of algorithms and protocols for MSN-related management issues [8], investigation efforts on how to effectively engineer MSN middleware solutions are still at a very initial stage. To offer a useful survey on MSN middleware architectures and supports, which starts filling the gap toward effective guidelines for MSN middleware engineering, we propose to analyze the space of possible design choices by considering existing MSN types and by identifying the variety of MSN structures enabled by current pervasive technologies (see Fig. 2).

First, we start with the distinction between infrastructure-based (typically centralized) and completely decentralized MSN solutions [8]. On the one hand, infrastructure-based MSN includes middleware solutions and applications that are typically based on centralized storage and processing of all application/context data; for instance, to this category belong all the applications where mobile users who join a new community transfer a copy of their current context metadata

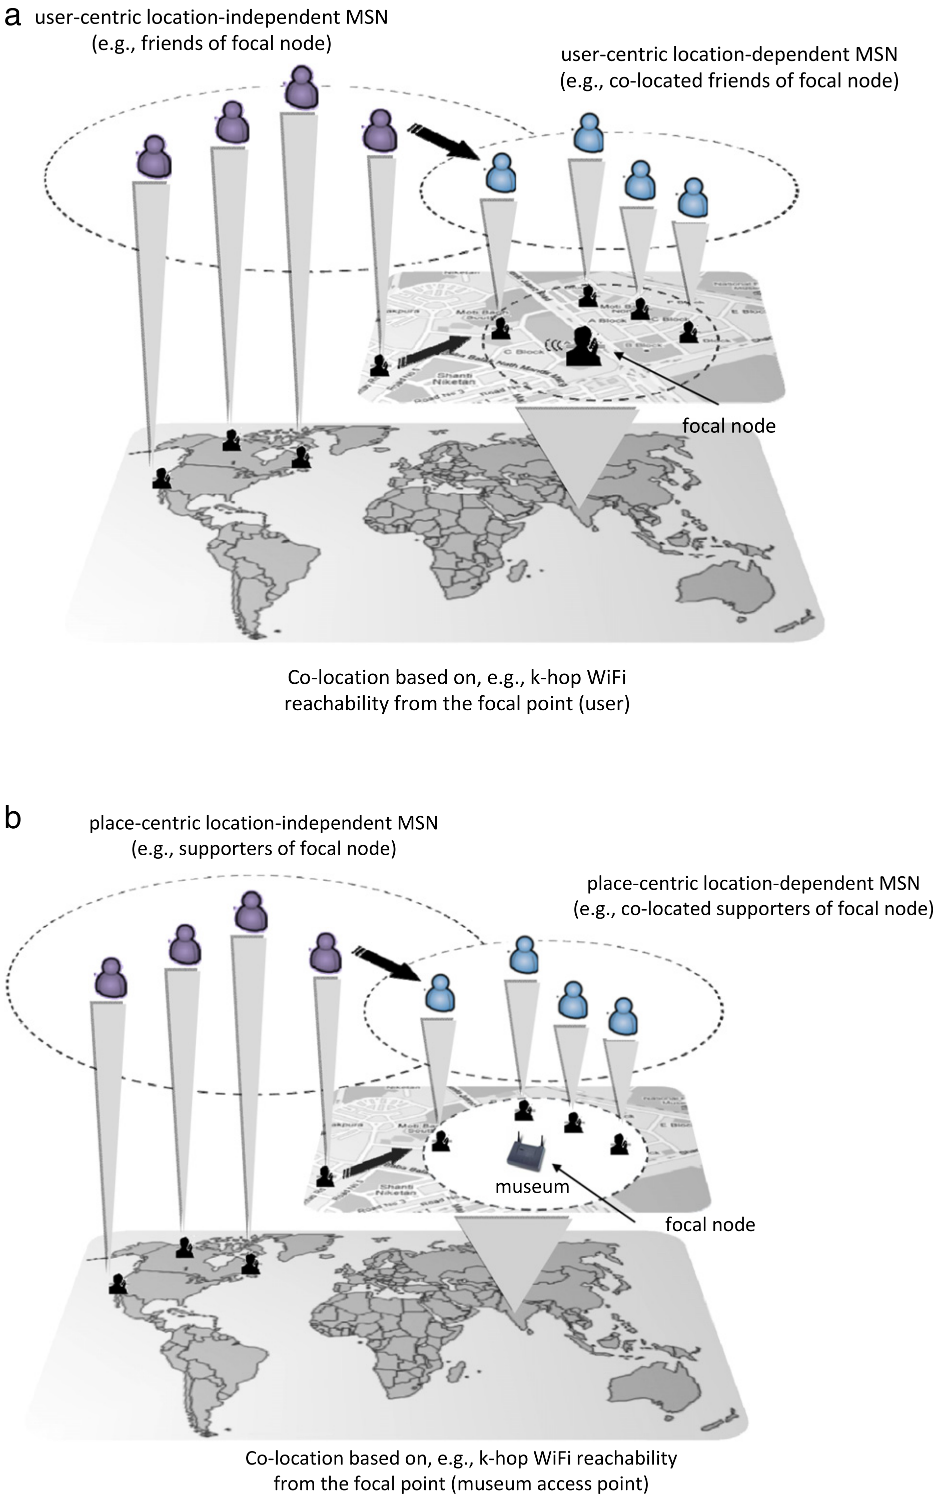

Fig. 2. A classification of MSN structures: (a) user-centric location-dependent/independent MSN, and (b) place-centric location-dependent/independent MSN.

(including social ones) to a server-based and logically centralized infrastructure; this centralized infrastructure is typically in charge of metadata processing/inference and community formation. On the other hand, decentralized MSN includes social networks only enabled by opportunistic peer-to-peer ad-hoc contacts among mobile users, e.g., who share and exchange application/context metadata only when in proximity.

Next, we extend this simple classification with additional dimensions that provide a more fine-grained overview of infrastructure-based and decentralized MSNs. As depicted in Fig. 2, we propose different possible structures for MSN views, with a *MSN view* defined hereinafter as a virtual community of users, along with their shared devices/services/resources, dynamically built depending on possibly different organizing criteria. MSN users are typically interested in a personalized MSN view of their community (and of the associated opportunities of group communications and resource sharing), in order to enable the visibility of only the opportunities of expected greater interest and to achieve scalability; personalized filtering and scalability are recognized as crucial factors, especially in mobile provisioning environments.

The rationale behind the schema in Fig. 2 is twofold. On the one hand, we propose this schema to provide middleware designers with various MSN view choices, which middleware solutions could integrate to achieve different tradeoffs in terms of performance, overhead, and "completeness" of social relationship inference. For instance, completely decentralized and local views may offer partial visibility of social ties, but with reduced connectivity needs and limited overhead. On the contrary, MSN views deriving from globally considering all the context metadata associated with all users potentially belonging to the community are likely to offer a more complete determination of a user's social ties, but for example may imply the transfer of a large amount of social-related metadata toward a server infrastructure. On the other hand, the proposed schema is general enough to address the need to support MSN applications with different degrees of social space virtualization, from completely virtualized to geographically-bounded and hybrid mobile communities [12].

Let us note that the MSN view taxonomy in Fig. 2 does not impose to middleware designers any constraints on the networking mechanisms/protocols/algorithms that can be used for inter-node communication and content sharing, as well as on the social network metrics [8] adopted to either identify the nodes belonging to a view or quantify the strength of node social ties; examples of such metrics include the similarity metric that groups nodes upon common interests [13] or the metric indicating social tie strength based on encounter frequency and duration [14]. Also note that other classifications are possible, but our approach permits a sufficiently general and system-oriented perspective, capable of distinguishing different design/implementation solutions with significantly different tradeoffs between development/runtime costs and flexibility, as better and more explicitly detailed in the following.

More in details, we propose the MSN structures depicted in Fig. 2 and based on two organizing criteria, considered as the two most relevant factors in many practical MSN applications: (i) the *physical location* of mobile users/devices, possibly with regards to specific places of interest, and (ii) the *focal node*, i.e., the entity around which the MSN view is built up to reflect the social ties of its major interest.

With regard to the first criterion, the physical location of a mobile user (and of her portable devices) or of a physical place, such as a stadium, a bar, a museum, can be used as key metadata to infer social relationships among mobile users and to create communities based on physical contacts and proximity (in addition to the more traditional consideration of common interests). In other words, physical location, together with the related concepts of mobility patterns and proximity, may be used or not to establish well-defined boundaries for MSN, thus generating the distinction between location-dependent and location-independent MSNs, as depicted in Fig. 2. In particular, in the case of location-dependent MSN views, we propose to consider the physical location of a special entity (called focal node and described in the following) as the center of a logical grouping abstraction, i.e., the MSN locality. The MSN locality is defined as the set of all mobile users whose devices are connected to the focal node by a route path of at most h hops. The h value can vary depending on the application scenarios and wireless networks employed. Mobile users within the same locality of the focal node, i.e., nearby users, may belong to the same location-dependent MSN, depending on the application-specific social network metric adopted [8]. On the opposite, location-independent MSN aggregates entities depending only on their context metadata (e.g., preference profiles, social interests, and affinities) excluding physical position.

The other organizing criterion in our taxonomy is the focal node, i.e., the entity determining the first-level aggregation principles based on which the MSN middleware builds its dynamic communities. In particular, we propose to categorize MSN solutions as *user-centric* or *place-centric*, depending on whether the focal node role is played by a user or a physical place. In *user-centric* MSN, the focal node is the user and the MSN view is built on the basis of interests, attributes, and/or needs of that user. To determine one user's view, user-centric MSN solutions can tie together users/resources not only independently of their geographic coordinates, as in traditional virtualized communities (called *location-independent user-centric* MSN), but also on the basis of their physical proximity (called *location-dependent user-centric* MSN). In particular, location-dependent MSN solutions promote opportunistic interactions between the ego user (focal node based on which the MSN view is built) and previously unknown nearby users, by exploiting context metadata either temporarily transferred to a MSN server infrastructure (infrastructure-based) or opportunistically exchanged among neighbors (decentralized). Location-dependent MSN views are typically characterized by a high degree of dynamicity, e.g., depending on user/device mobility and consequent frequent modifications of resources in the view. This has a non-negligible impact on achievable performance: there is the need of properly trading between prompt reactivity to changes and coordination overhead. Generally, better scalability could be achieved whenever it is possible to strongly limit the number of participants to the MSN view (limited scope of interest).

In contrast, in *place-centric* MSN, the focal node is a physical place that acts as the space where social encounters/interactions occur, i.e., acting as the aggregation criteria to put together participating users and their resources into a MSN view. In this case, the view is built on the basis of place attributes and/or physical co-location of users within the place boundary. In this MSN class, given that the focal node is the place, a rich set of mechanisms and facilities are available to define boundaries around a place: for instance, they may be specified by an n-sided polygon where each point is represented by a set of n GPS coordinates, or a circle where the GPS coordinates of the place represents the center. Place-centric MSN solutions can also be classified into location-independent and location-dependent ones. On the one hand, in location-independent but place-centric MSN, a MSN view groups together users, regardless of their current physical locations, who have exhibited their strong interest for a given place, e.g., fans of the Pitti Palace museum in Florence or followers of the Café Florian in Venice. In place-centric location-independent MSN, the place tends to aggregate and have global knowledge about the community members interested in it, independently of the fact that these members are currently located there. The benefit of such a global knowledge is easier dissemination of information, such as special event advertising, under the full management, business, and security control of the place. Let us note that, in principle, other types of users' interests (not only in a physical place) could be used as focal nodes in a MSN classification; we decided to consider the place as a special case because in place-centric MSN implementations, both location-dependent and independent, context/social metadata are frequently hosted on MSN servers belonging to the place, e.g., the MSN server infrastructure of a museum, thus suggesting different engineering design guidelines for effective MSN provisioning.

On the other hand, a place-centric location-dependent MSN view groups mobile users currently located within the place boundaries. The geographic boundaries of the place define a locality and delimit the scope of user/resource discovery in order to form the applicable MSN view. Those MSN solutions can be exploited to promote more focused and social-aware fruition of the services offered by the currently visited place, as usual for tourism-oriented MSN applications. About locality-based scalability advantages, we can apply similar considerations to the ones made above for the user-centric case.

It is worth noticing that, in both user- and place-centric MSN views, the resource/service discovery scope in location-independent MSN is determined only by social interests and affinities (and can be potentially boundary-less), while the discovery scope of location-dependent MSN depends on proximity with the focal node. In addition, let us clarify the motivations why we claim that it is relevant to keep distinct location-independent and location-dependent MSN views, and the benefits deriving from this distinction.[2] One primary advantage is that any MSN mobile user can bring with her a set of social ties resulting from her location-independent MSN view. When joining a location-dependent MSN view, the user can act as a bridge between different virtual/physical communities and decide to bring in (partially or completely) her social metadata related to her location-independent MSN view. This allows the user to share her social ties with other opportunistically encountered users, e.g., belonging to the same location-dependent place-centric MSN view, thus possibly establishing new, either temporary or time-lasting, social relationships that were previously unforeseen. Therefore, location-independent and location-dependent MSN views may have intersections, also only temporarily. For instance, let us consider the case of a young student (Alice) traveling from her college town to the closest city to visit a large bookstore. Alice may be the user of a location-independent MSN application where she has a direct social tie with Nick, one of her college friends. While having a coffee break in the bookstore, it may turn out that Nick also currently belongs to Alice's view in another location-dependent MSN application because he is in her proximity (Nick is visiting the city on the same day). Alice could also establish a new social tie with a previously unknown person, i.e., Bob, currently co-located in the bookstore. Then Bob becomes a member of Alice's location-dependent MSN view and eventually Alice can decide whether to prolong this social tie beyond her staying at the bookstore. By temporarily/permanently sharing context metadata (including social relationships) with Bob, Alice can play the role of bridge between Alice's and Bob's social communities; thus, the MSN views of the two users (and potentially of all the other users belonging to the two MSN views) may be temporarily/permanently widened; to achieve this goal, different tradeoffs between communication overhead, scalability, and completeness of the inferable social ties of interest, may be adopted depending also on using either infrastructure-based or decentralized MSN solutions.

2.2. Spatial and temporal scopes

We claim the relevance of introducing two additional dimensions in our taxonomy. These dimensions relate to: (i) the scope of the collection of context metadata used to determine the MSN view (*extended spatial scope*, global vs. partial); and (ii) the ability to store the time evolution of those metadata (*extended temporal scope*, history-based vs. instantaneous).

Let us rapidly discuss and exemplify how spatial and temporal scopes can affect social metadata analysis, the dynamic creation of social communities, and middleware performance/scalability. On the one hand, with global spatial + temporal scope, the MSN middleware may exploit the full visibility of context metadata, e.g., to infer additional social relationships or detect emerging global behaviors and their evolution in time. Global spatial + temporal scope is usually at the basis of techniques for so-called detected social networks [8,15], i.e., communities that are determined by analyzing context such as encounter/re-encounter traces, contact frequency, etc. Detected social networks usually call for MSN supports that logically centralize the storage of context metadata for all users and places, typically accomplished in infrastructure-based MSN solutions (see Fig. 3). For instance, by focusing on user-centric MSN, middleware can offer support for possibly complex and time-consuming offline processing algorithms to compute user-centric social affinities, such as user groups of interests, preferred activities with friends, and preferred visited places on a regular basis. Similarly, place-centric MSN with global scope can provide information on how crowded a place is at different times during the day, what kind of users visit that place, and which social events are attracting more attention.

[2] It could be argued that location-independent MSN views map to traditional OSN communities, whereas only location-dependent MSN views are peculiar and specific of MSN scenarios. Our distinction reflects our aim of having MSN structures that enlarge the possibilities of creating and reasoning about social ties and that do not to replace but rather complement existing OSN communities.

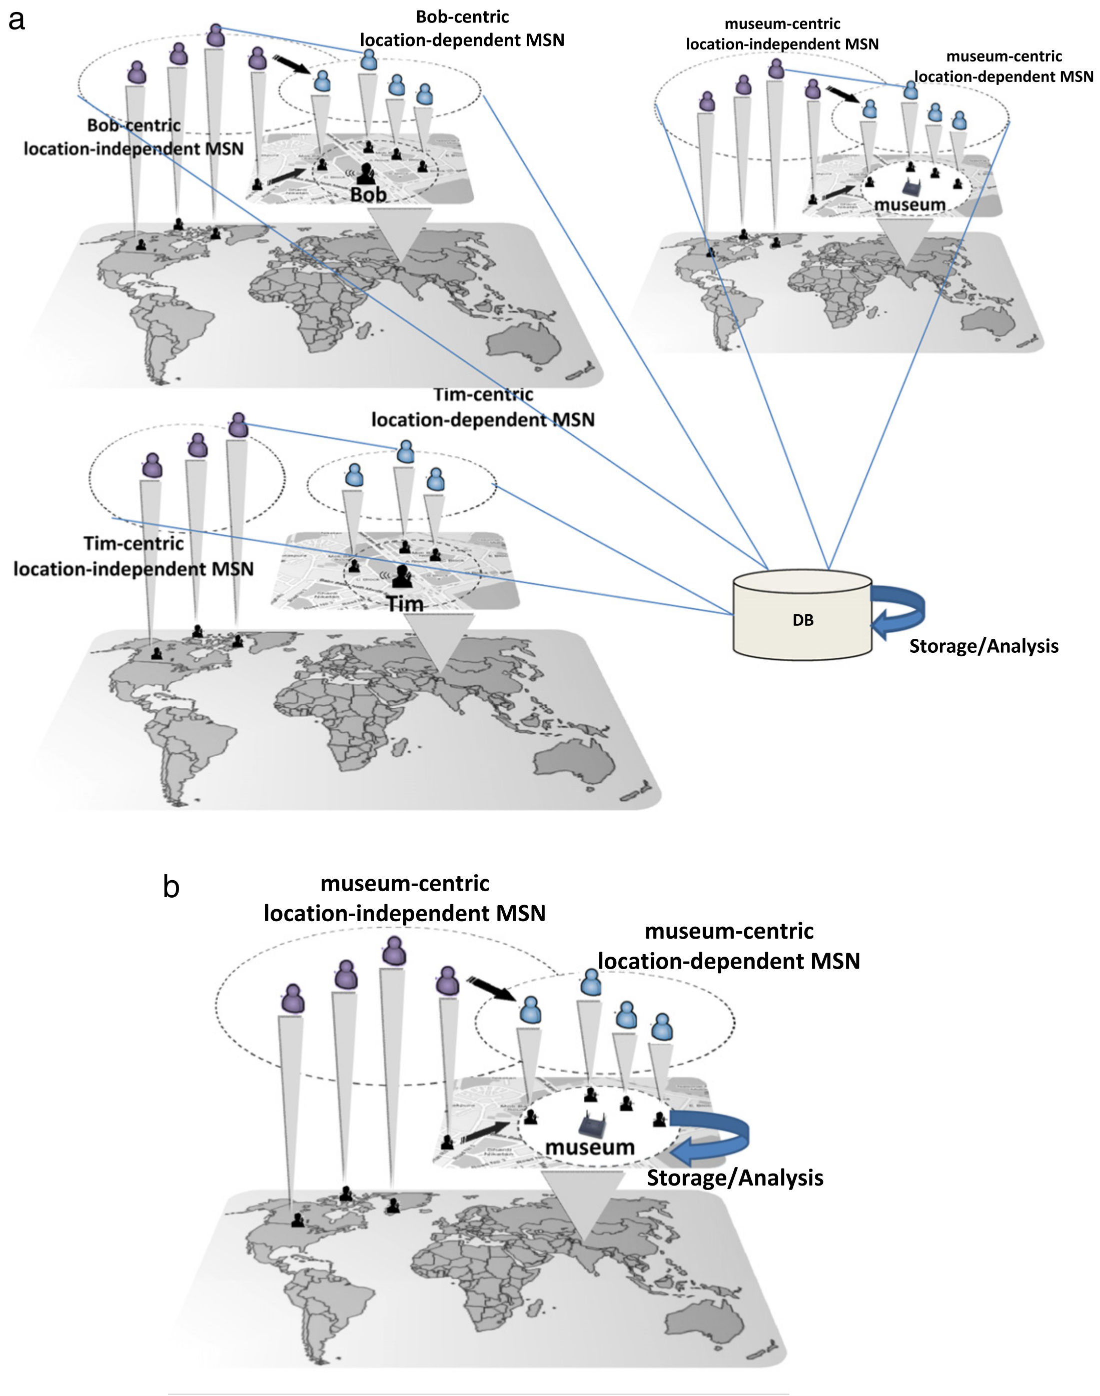

Fig. 3. (a) Global and (b) local spatial scopes. Both cases may be with either history-based or instantaneous temporal scope.

On the other hand, context metadata could be more suitably provided to the MSN supports only locally and for a limited time interval, possibly in a partially anonymized way, such as when self-reported[3] social network approaches are adopted, or when only decentralized MSN architectures are possible, or when privacy requirements are strict. Sometimes the choice of partial scope for metadata visibility/storage (e.g., MSN middleware that organizes disjoint domains with only local visibility

[3] The *self-reported community* term is used in the literature to indicate a social network whose formation is determined by the social relationships metadata explicitly provided by one participating entity, e.g., the user focal node in case of a user-centric approach. For instance, Alice's self-reported community may consist of all the users belonging to the Facebook social relationships graph provided by Alice to the MSN application/middleware.

of visiting user metadata) can help in achieving scalability over large deployment environments. Moreover, let us note that the choice of partial spatial + temporal scopes can heavily influence the business model associated with MSN application provisioning, e.g., by preventing the concentration of huge amounts of context metadata to a single provider and server infrastructure, or by opening interesting opportunities for fully distributed metadata ownership.

To clarify with a practical example and to better point out the usefulness of the proposed categories for MSN structures, let us consider the case of a family (with two parents and their little kid) and again the young student Alice, visiting a crowded museum. The family, Alice, and the museum authority can exploit different types of user/place-centric MSN middleware/applications. While in the queue, Alice can exploit her user-centric location-dependent MSN view to exchange art ideas, opinions, and itineraries with users with similar art affinities and possibly waiting in the same queue. Similarly, the museum authority can decide to entertain users in the queue with an art-related game that entitles a final free entrance award (place-centric location-dependent MSN). If the museum authority is willing to organize another game, always entitling free entrance awards, but opened to all users who manifested interest in the museum, it can suitably exploit a place-centric but location-independent MSN view. Depending on the type of scope, different queries and inferences can be carried out by an associated MSN middleware. For instance, if it adopts global spatial + temporal scope, it is possible to perform full analysis on collected context metadata, by analyzing the gender or age of people exchanging ideas while in the same queue or more generally visiting the museum.

3. State-of-the-art research efforts in MSN

As mentioned earlier, the surge of social networks, along with the widespread diffusion of smart phones, has recently stimulated relevant research efforts from a rather wide range of viewpoints [1,2,5–7,16]. There is a large body of proposals related to the integration of social networks and pervasive computing. The vision of people-centric sensing represents a significant research area along this integration direction [1,16]. The availability of a plethora of sensors onboard mobile devices, such as sensors for activity recognition, GPS for positioning, high-quality cameras, microphones, and accelerometers, enables a myriad of new sensor-based applications. In people-centric approaches, humans become the focal point of sensing: personal information can be monitored and archived (personal sensing), can be shared within social and special interest groups (social sensing) or with everyone for the public good, such as entertainment or community action (public sensing). People-centric sensing sits at the nexus of several research disciplines, including sensor networking, pervasive computing, machine learning, and social networking. Based on data collected from a multitude of sensors embedded within a smart phone, it is possible to infer the status/mood of users and their current activities.

A strong link also exists between opportunistic and social networking [2,3,9]. In opportunistic networking, context (and social) metadata, used at different levels of the networking stack, play a central role. In general, opportunistic network topologies and connectivity are quite unstable and unreliable. When traditional information about network status cannot be exploited, e.g., to find the available paths from source to destination, context-based prediction of contact opportunities between nodes allows more effective message delivery [6,7]. Some proposals have emerged that highlight the importance of context- and social-awareness in the design and optimization of protocols for opportunistic networking, especially of routing and content dissemination protocols [6–8]. For example, in [7] several social-aware policies have been defined and tested to evaluate when social awareness allows achieving better performance and effectiveness in data dissemination. In addition, a context- and social-aware middleware for opportunistic networking has been proposed that automatically learns context and social metadata and that uses this information in order to predict users' future movements [6].

MSN is the other relevant research area that strongly witnesses the integration of social networks and pervasive computing. Relevant research work has been recently accomplished on the theme of mobile virtual communities along different perspectives [12] and in different application domains, e.g., for tourism journeys, collaborative student environments, mobile sharing of music, and mobile business [17,18]. From the technical perspective, research in the field of mobile virtual communities mainly involves the design and development of effective supports for community creation and management. Other related research areas include MSN application/resource deployment according to user expectations and investigation of emerging user behaviors within mobile social spaces. From a business perspective, research on mobile virtual communities concerns innovative business models to make profitable the provisioning of MSN applications. Let us note that the previous areas of people-centric sensing and opportunistic networking can contribute to provide enhanced MSN applications by offering relevant concepts and solutions that can be fruitfully integrated in MSN middleware supports. Being MSN middleware solutions the focus of this paper, we hereinafter restrict the scope of the related work analysis specifically and only on this topic.

3.1. MSN middleware proposals

Most of the initial MSN solutions considered only mobile devices as tools to access existing social networks. Recently, in addition to simply making pre-existing social networks available to mobile users, the need for building social networks spontaneously while in mobility has increasingly emerged. Along this direction several social applications on mobile devices have been developed, aiming at going beyond mobile-accessible versions of Facebook or MySpace.

One category of such applications allows users to find common interests among nearby people by using spontaneous ad-hoc and single-hop communications. Social Net [19], Social Serendipity [20], IYOUIT [21], and HIC [22] are examples of such category of MSN applications. Specifically, Social Net [19] infers shared interests between opportunistically encountered people (single-hop proximity for a sufficient time interval) by storing unique identifiers of nearby devices, by recording the time and duration of encounters, and by analyzing the patterns of physical proximity collected over time. For example, if two people regularly exercise in the same fitness center at the same times of the day, Social Net detects this pattern and assumes these two users could be interested in some forms of resource sharing. Social Serendipity [20] is another similar MSN application for Bluetooth-enabled smart phones: it maintains and offers a database of user profiles to recommend face-to-face interactions between nearby users who share common preferences. IYOUIT [21] is yet another MSN application mostly acting as a social context provisioning and sharing platform. Here context is typically centered on places/people that IYOUIT users can visit/meet. Finally, HIC [22] is a Near Field Communication-based MSN application that allows people to create friendship relationships on the spot when they meet each other. The HIC focus is on determining which context information is relevant for the action of friend connection and on how context data should be taken into account for application design. Another interesting solution is the WhozThat [23] protocol, which can enrich the previous applications with the possibility to bind nearby users with social network IDs and to import their social context from online social network applications into each user's local context.

Several existing MSN applications are designed from scratch by embedding the MSN management functionality into the application logic and by providing application-specific representation models for context and social metadata. These applications are still more proof-of-concept prototypes than industrial MSN applications supported by comprehensive frameworks to facilitate and support their cost-effective design, development, and deployment. It starts to be recognized that the availability of suitable MSN middleware can significantly improve the productivity of developers as well as the performance/features of MSN applications. Only to mention one central aspect, designers could take advantage of the same middleware support in different social computing applications, thus enabling rapid prototyping and facilitating MSN application interworking. Some authors [24] have recognized the need for externalizing social management functionalities without embedding them directly into applications; more in general, the rising agreement on the principle of MSN middleware suitability is pushing the realization of several support solutions [25–31]. Each of these middleware proposals supports different MSN structures and sets of facilities, with the deriving strengths and weaknesses. Given the existing plurality of approaches, we claim that an overall and structured classification of them can help in identifying common directions of solution and facilities (possibly already available for integration), as well as in determining primary technical challenges not yet addressed (at least efficiently and effectively). In the following, for the sake of comprehensive comparison and to clearly point out how the proposed taxonomy can help, we have decided to survey only the subset of most complete MSN middleware solutions, which provide similar sets of support facilities/functionality for MSN management.

3.1.1. MobiSoC

The MobiSoC middleware [25] provides a common platform for capturing, managing, and sharing the social state of communities. The community concept is a logical abstraction used to aggregate users and physical places. It defines the boundary for collecting user/place characteristics, location information about the users belonging to the community, and user/place socially-relevant data. The community social state consists of people profiles, place profiles, people-to-people affinities, and people-to-places affinities. In particular, MobiSoc allows the potential creation of both user-centric and place-centric MSN views, by realizing an extended spatial scope with a centralized architecture. The MobiSoc social management model relies on a centralized server that globally collects people/place profiles and location metadata. It exploits learning algorithms over globally collected information to infer emergent geo-social patterns, such as social ties between people, and associations between people and places. Such a MSN approach with global spatial scope simplifies geo-social pattern identification and privacy policy enforcement over sensitive data. In addition, MobiSoc provides history-based MSN structures (extended temporal scope). The social state evolves continuously over time as new user profiles, social ties, place-related metadata, and events are created. The MobiSoc guideline is to maintain, as far as possible given the storage and processing time constraints, the whole history of social state evolution. Also user mobility traces are kept over a long time interval to enable complex geo-social pattern recognition.

By delving into finer details, MobiSoc provides specific modules for collecting users, places, and location data. (i) The *people module* for inserting/modifying/acquiring users' identification data, social interests/ratings, and preferences; (ii) the *place module* for supporting the collection of geographical data and maps for buildings, offices, and outdoor locations, as well as for introducing and modifying social events associated with a place; and (iii) the location module for receiving and storing location updates from mobile devices. MobiSoc exploits XML-based user/place profiles. In addition, it provides several modules for social state learning, i.e., for discovering and representing relationships between persons and places. Learning algorithms are adopted to determine people-to-people and people-to-place affinities. For instance, MobiSoc allows finding out that a user attends research seminars regularly, plays tennis every Friday, or works together with another user. It is also possible to find out how crowded a place is at different times in the day, popular social events which happen at that place, or the demographics data of people who visit the place frequently. Moreover, based on the social affinities computed from factors such as similar interests, similar backgrounds, common friends, or common places, MobiSoc automatically works on the discovery of previously unknown emergent geo-social patterns.

In particular, to compute people-to-people affinities, MobiSoc examines all user profiles in the community and analyzes every pair of user profiles and user mobility traces along several social factors (common events, mutual friends, etc.), as well as geo-social factors, such as co-presence or presence in a given place in different time intervals. To identify associations between people and places, MobiSoc adopts a learning algorithm that tries to detect previously unknown ad-hoc groups and their associated meeting places. The algorithm performs time-based clustering on users' mobility traces to infer repeated user co-presence at the same place. Users with a higher Degree of Co-Presence (DCP) are likely to belong to the same ad-hoc group. Then, to finally determine whether a place is a group hangout, the algorithm checks the users' DCP with the total number of visits. The required DCP is chosen on the basis of a theoretical analysis that allows balancing the tradeoff between group member identification and false positive percentages. It is worth noticing that MobiSoc also provides an event manager facility that allows MobiSoc-based applications to receive notifications when parts of the social state in their view change, such as a new match based on newly identified social affinities.

A key feature of MobiSoc is also its simple solution for privacy management and enforcement. Its security model relies on a centralized trusted entity that works to enforce privacy uniformly in any deployment environment: users provide their privacy policies to a trusted centralized entity, which enforces them on their behalf. The negative side effect of this simplicity is that the trusted entity controls and maintains all the community social state, with the related privacy-sensitive potential issues stemming from easy possibility to cross-correlate all the available metadata.

3.1.2. SAMOA

The SAMOA middleware [26] supports the creation of anytime anywhere MSNs. In particular, it allows mobile users create roaming social networks that, by following user movements, provide at each instant the view of all possible neighbors of expected interest. The SAMOA roaming social network is centered on a user (the so called ego-user) and is created based on the visibility of both current user's physical place and her requirements/characteristics. The ego-user is the manager of her own social network; she defines the criteria guiding her MSN view (discovery scope). The place abstraction is used to determine the discovery space of a roaming social network and to limit it. In the SAMOA terminology, the place is the set of SAMOA users who are physically proximate to the ego-user and are eligible to become members of the ego-user's social network. SAMOA follows a decentralized MSN approach with partially scoped location-dependent views. The focal node can be either a mobile user or a physical place, thus supporting both user-centric and place-centric MSN views. In the latter case, SAMOA can be configured to follow a place-centric approach by simply forcing a physical place to play the role of a fictitious ego-user. In addition, mobile users acting as ego-users can exploit two different types of social networks: (i) *place-dependent MSN* that keeps track of only the members currently co-located with the ego-user, and (ii) *global MSN* that persistently records the whole set of place-dependent views created over time as the ego-user moves across places (extended temporal scope). On the one hand, place-dependent social networks let managers easily discover co-located users of interest when they want to establish one-shot and transient interactions. On the other hand, global social networks let managers create application-dependent interaction histories, which can enable more complex collaboration strategies and patterns.

The SAMOA middleware integrates a set of common management facilities for semantic-based MSN extraction and management, such as tools for specifying and checking profile correctness, social matchmaking algorithms, and services for creating and maintaining place-dependent and global MSNs. SAMOA represents user/place profiling metadata in the Ontology Web Language (OWL) to promote interoperability. It includes OWL-specific semantic matching algorithms that exploit reasoning based on subsumption and description logic for analyzing profiles and inferring social semantic compatibility. SAMOA also provides its MSN application developers with other management facilities to support naming, detection of co-located SAMOA entities, and communication. Its security support is instead still in its infancy and under development.

3.1.3. MobiClique

MobiClique [27] is a MSN middleware designed to operate on mobile ad-hoc network settings only, where Bluetooth-enabled devices communicate directly (with no need for additional network infrastructure) as they meet opportunistically. This middleware allows mobile users to maintain and extend their online social networks through opportunistic encounters, thus promoting user-centric MSNs. Hence MobiClique does not depend at all on centralized servers but only relies on opportunistic connections between neighbors (local spatial scope). When proximate users meet, if their profiles share a pre-defined relationship or interest, they can decide to become "friends" or to exchange some kinds of content. When two users become "friends", their social graphs can also take the opportunity to expand, because they discover previously unknown users included in the social graph of the encountered peer.

MobiClique associates each mobile device with a social profile consisting of a unique identifier, some personal information, a list of friends, and social groups of interest. It adopts a simplified profile format, which is not semantic-based, just for the sake of fast prototyping; the proponents claim that this format will be soon replaced by a richer and standard profile format, such as Friend of a Friend (FoaF), Social Graph API, or OpenSocial [32,33]. The locally maintained social profile of a mobile device is initialized and kept synchronized with an existing online service for traditional social networking, such as Facebook. The assumption is that, even though mobile devices mainly operate in an ad-hoc mode, they can occasionally exploit Internet connectivity. About security, when two mobile devices meet, they perform user identification/authentication based on public-key cryptography. Once authentication is successfully completed, the two devices can exchange their full social profiles and/or shared content. To support opportunistic social encounters, MobiClique provides a set of functionalities, including neighborhood identification based on Bluetooth device discovery or broadcast beacons on a well-defined WiFi Service Set IDentifier (SSID).

Table 1
Key features of existing MSN middleware solutions.

MSN middleware solutions	User-centric vs. place-centric MSN	Extended vs. local spatial scope	History-based vs. instantaneous MSNs	Mechanisms for social ties	Security support
MobiSoc [25]	User- and place-centric	Extended	History-based	Learning algorithm	Privacy support
SAMOA [26]	User- and place-centric	Local	History-based and instantaneous	Semantic inference	No security
MobiClique [27]	User-centric	Local	Instantaneous	Syntactic pattern-matching	Entity's authentication support
Yarta [28]	User-centric	Local	Instantaneous	Semantic inference	Access control support

3.1.4. Yarta

Yarta [28] is one of the most recent and complete MSN middleware proposals. It can execute on both fixed computers and smart phones, allowing users to create user-centric MSN and to maintain their MSN metadata locally (onboard of their devices). In particular, Yarta is designed by taking into careful account the heterogeneity of typical metadata available for inferring geo-social relationships and the intrinsic decentralized nature of MSN applications. As a key feature, Yarta provides an expressive and flexible model to represent mobile social ecosystems, defined as the rich set of interactions occurring between MSN participants. Yarta proposes a set of original first-class entities (building blocks for the creation of mobile social ecosystems) and relationships between these entities, which are considered sufficient and necessary to design any type of MSN application. More specifically, the Yarta model defines the following first-class entities: (i) *agent*—an entity having beliefs, desires, and intentions according to the Belief-Desire-Intention design model for intelligent agents, (ii) *place*—a physical entity located in space, (iii) *content*—an entity that conveys an information, possibly to be shared, such as a picture, (iv) *topic*—a subject of interest for MSN users, and (v) *event*—an entity whose existence is characterized in a three-dimensional space and time environment. Several first-class binary relationships might hold between these entities; Yarta defines and provides them to application designers with no additional effort.

By delving into finer technical and implementation details, a Yarta MSN ecosystem is represented as a graph of Resource Description Framework (RDF) subject-predicate-object triples, with each statement describing an attribute and its value. The set of RDF statements defining social information is linked as a graph of nodes and arcs. The choice of RDF has shown to be a good tradeoff between simplicity of metadata representation and possibility of enabling some forms of reasoning via the association of formal semantics with the metadata model. Reasoning capabilities are key enablers for social metadata interoperability between Yarta users. Each user owns a personal knowledge base that includes all its MSN metadata; she can decide to share her knowledge base with other users by exploiting the Yarta metadata model for interoperable representation. Because the Yarta model is based on a graph representation, it is well suited to support the retrieval and representation of partial mobile social ecosystem knowledge, since graphs can be split and merged based on users' application needs.

Yarta provides several management functionalities to support the design of MSN applications, including services for ecosystem creation/update and for social metadata extraction from multiple sources, thus allowing developers to collect and reuse metadata from other existing and more traditional social networks. In addition, from the security support view point, Yarta only focuses on solutions for advanced access control to the knowledge bases that are shared by participating users.

3.1.5. Comparison of existing MSN middleware

As it stems from the concise synoptic comparison in Table 1, most MSN middleware proposals nowadays adopt a user-centric approach with instantaneous MSN views that are stored in a local partitioned way (local spatial scope). MobiSoc is the only middleware proposal that provides global and history-based social state with a centralized management model. The choice of local vs. global and instantaneous vs. history-based views has relevant implications on social metadata availability and simplicity of analysis, as well as on middleware architectures and types of social communities that can be formed. The MSN view partitioning enables local immediate availability of social metadata on mobile devices (possibly with partial and temporary inconsistencies if compared with the real global state), even in the case of intermittent connectivity to any centralized server. Local MSN views can rely on decentralized middleware with local and autonomous point of management onboard of the mobile devices themselves. On the contrary, centralized architectures provide effective support to global social state maintenance. It is also worth noticing that identifying emergent geo-social patterns is harder to achieve with MSN metadata partitioned into different localities and stored at different participants. For instance, because of its characteristics and facilities, MobiSoc allows identification of emergent community patterns, thus enabling the possibility to form detected social networks, in addition to self-reported ones [8]. On the contrary, the other middleware proposals in Table 1 do not allow inferring emergent geo-social patterns.

All solutions tend to provide some basic set of middleware functionalities for collecting location metadata, for discovering nearby devices, and for exchanging and reasoning over social metadata. However, there are several central issues still open. The first and most evident aspect is still the lack of a standardized general architecture that clearly defines the needed support facilities that a MSN middleware should offer. For instance, the middleware proposals in Table 1 highly differ

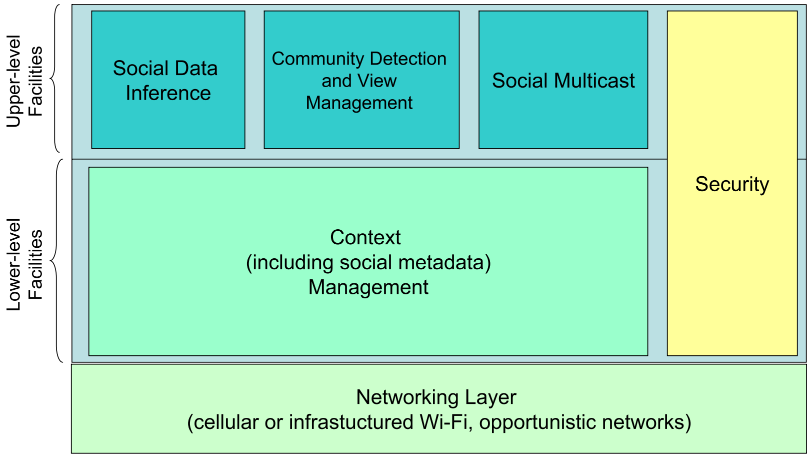

Fig. 4. A conceptual reference architecture for MSN middleware solutions.

in terms of the supported facilities. In addition, they all lack integrated support for social-aware communication and full context awareness. Explicit and clear APIs for social communication should be provided; moreover, full context-awareness support should be incorporated to enrich social contacts and ties. The previously described Whozthat system is an example of integration between social and context-aware networking, but MSN middleware solutions along this direction are still lacking. Even though context characterization, modeling, and reasoning have been extensively investigated in the pervasive computing literature [34–37], there is still the need for a general context management support as a common feature that a MSN middleware could integrate, in particular by giving the opportunity of configuring it dynamically as required by running applications. In addition, social context has not been incorporated in more traditional context middleware; only some first efforts toward the integration of social- and context-awareness have recently emerged, mainly in the fields of opportunistic networking and novel context/social management frameworks for wide-scale ubiquitous computing [6,24,38].

Dynamic middleware configurability is another relevant technical challenge that is inadequately addressed by the current literature. Middleware designers should be able to integrate different protocols, algorithms, and solutions into the same MSN middleware, so to enable the proper middleware configuration and solution selection only at deployment/run time depending on application requirements, mobile device settings, and available network infrastructures. Finally, other relevant technical issues still have to be addressed to favor commercial and practical feasibility of MSN over wide-scale deployment environments, namely MSN-specific solutions for security, interoperability, and scalability.

4. MSN middleware: architecture and primary design guidelines

To design a comprehensive MSN middleware, as it stems from the comparative analysis of state-of-the-art solutions in Section 3, several hard management challenges need to be addressed, ranging from community detection in the presence of node mobility to social relationships inference, from social-aware group communication to MSN view management, security, and interoperability. Some of them have already received initial attention in the existing literature [8], with various solutions already emerging; others call for additional research work and especially wide experimentation over industrial deployment environments. The presentation of the specific protocols to solve these management issues is not the goal of this article. Instead, based on our state-of-the-art analysis, here we aim at identifying and discussing the key components of a general MSN middleware by which it can be architected. In fact, MSN application development can be greatly simplified if programmers can rely on ready-to-use, horizontal, and application-agnostic middleware facilities to address the above management challenges. In particular, we propose a conceptual middleware architecture including facilities organized in two logical layers, built on top of the networking support (see Fig. 4). The upper-level facilities offer support to discover social ties of interest (*social data inference*), to create communities (MSN views) depending on discovered social ties and to manage MSN views (*community detection and view management*), to allow social-aware communication among MSN view members (*social multicast*). The basic level includes lower-level functionality to collect/manage context metadata and to support location tracking (*context management*). At both levels, crossing them, is the *security service*: at the upper level, it provides the application with the possibility to define and register security preferences/policies for its users, for instance, who can access and navigate a user's social graph, when, and how; instead, at the lower level, it offers all the mechanisms needed to address confidentiality, integrity, authenticity, and to enforce access control/privacy according to user's preferences. The MSN middleware facilities are layered on top of a networking support, which may be specific for different technologies and deployment environments, from more traditional networking infrastructures such as cellular and infrastructured WiFi, to opportunistic networking that enables routing/dissemination with different context/social visibility degrees (context-oblivious, partially context-aware and fully context + social-aware) [6,7].

More in details, the *social data inference* facility allows programmers to identify complex social relationships by combining different types of context. This facility may exploit context metadata coming from different sources, e.g., from the underlying sensor management facility and from OSN social graph repositories. To infer complex and rich social relationships (in addition to the ones explicitly offered by the underlying sources), it can largely benefit from global visibility of context metadata. It can exploit different social metrics to build social ties, thus potentially providing different social graphs up to the application level. The choice of the most suitable metric to adopt usually depends on application requirements and deployment environment characteristics. On the basis of the adopted social metric, the social data inference facility can properly tag social graph nodes with weights, such as in terms of centrality, closeness, connectivity degree, social interest, . . . , by producing a directed weighted social graph. Various techniques are available in the literature, as it also emerges from the state-of-the-art analysis in Section 3, ranging from semantic reasoning to machine learning. Semantic reasoning is more suitably adopted if context is modeled according to ontology-based representation; machine learning techniques can be effectively exploited to find recurring patterns, such as co-presence or re-encounters of mobile users within the same physical place, and from these recurring patterns to extract meaningful social relationships among users.

The *community detection and view management* facility allows developers to seamlessly determine and manage social-relevant groups of MSN users. In particular, it allows the formation of social communities according to both the major followed approaches, i.e., self-reported or detected social networks. To this purpose, the facility navigates the inferred, directed, and weighted social graphs to extract meaningful portions of them. In other words, the community detection facility performs a filtering reduction/prioritization of the parts of the social graphs to (primarily) take into account according to application-specific requirements and the needed performance tradeoff. The filtered social graphs map to the MSN structures of our taxonomy in Section 2. This facility can integrate protocols, metrics, and algorithms proposed for community detection [8]. This facility is also responsible for managing the created MSN views. In particular, it provides MSN developers with the possibility to access, query, and navigate MSN views. In addition, in the case of history-based MSN, the facility is responsible to manage the time evolution of MSN views by tightly coordinating with the social data inference at the upper level and the context management at the lower level. The community detection and view management facility should be notified of any community-relevant change over time on context metadata, so as to trigger a re-calculation of social graphs and update the MSN views accordingly.

The *social multicast* facility provides socially-aware multicast support. The service not only allows group communications among the members of a MSN view in a seamless way, but also enables easy notifications about relevant social changes/events by coordinating with the underlying context management layer. For instance, the facility can be configured to notify a registered application of the co-location of two users, of a user's presence at a given place, or of a new social match based on newly identified social affinities.

All upper level middleware functionality is based on the *context management* facility in charge of context acquisition, representation, and delivery. Several existing context frameworks can be integrated and exploited to support some basic facilities for traditional context management [39,40], whereas only few existing middleware platforms consider social metadata or social community abstractions for context management [6,38]. Through coordination with upper level facilities, the most suitable scope of acquired and delivered context metadata is determined: for instance, context management can be configured to immediately deliver notifications about status changes of either any MSN view member or only of the focal node user (other status changes are perceived only after explicit polling).

4.1. Primary guidelines for architectural and design choices

To provide a general-purpose, horizontal, and effective MSN middleware, many technical challenges need to be addressed from the perspective of both adopted protocols and design/deployment choices; several of them are still very open as shown by the recent related literature [8]. The detailed presentation of emerging state-of-the-art protocols for MSN support, e.g., for enhanced forms of community detection, view management, social data inference, or social multicast, is not the goal of this article, also because some first good surveys about that have recently started to appear [8]. Instead, by using our taxonomy proposal, here we only and originally focus, in a concise way, on some main technical challenges whose requirements may have a relevant impact on effective MSN middleware architectural and design choices, especially depending on the type of MSN middleware that fits a specific deployment environment or application domain.

Standardization of MSN interfaces, e.g., for exchanging social graphs and socially-aware data sharing, is crucial when different MSN applications need to interact. This is particularly important because of the dynamic and unpredictable nature of users' encounters in open deployment environments. It is possible to address MSN interoperability issues by working along several directions. First, MSN requires expressive and interoperable representation models to describe users and places, as well as the different relationships stemming from their interaction, e.g., event participation, percentage of crowd at a place in a given instant of time, or user's current activities in a place. A common trend is to exploit ontology-based data models and semantic technologies to represent social graphs and metadata [41]. In particular, complex and formal representation models, such as OWL-based ones, are demonstrating to better fit MSN scenarios with global spatial scope and non-real-time (e.g., offline and a-posteriori) context analysis to establish new social relationships and alike. More lightweight solutions enabling only simplified forms of metadata openness and understandability, based on less formal representations such as RDF, are more frequently used in local spatial scopes where context is exploited locally, typically by running (partial) context processing at one user's node (e.g., her smartphone). Of course, interoperable representation is also important to allow users' mobile devices interacting with their surrounding smart spaces and the wide spectrum of sensors possibly available there. Let us note that some smart space management frameworks are considering the opportunity of open and interoperable sharing of sensor data to enable technology/vendor-agnostic interoperability and to reduce integration costs. Also in this case, solutions are emerging in terms of open data models, semantic technologies, and lightweight representations of social relationships [42]. Moreover, it is central to push the standardization process of MSN middleware APIs to leverage the market of interoperable MSN applications, e.g., crucial to make feasible and cost-effective the realization of a place-centric MSN application to a museum owner: a few initiatives on open social APIs exist [33], but they are still considered immature in practice, mainly because of too rigid dependencies on the owners of associated OSN applications. In addition, still open technical issues associated with context interoperability relate to privacy and access control: MSN applications generally handle sensitive information per se, which can also be used to infer additional metadata, sensitive in their turn. Privacy implications further increase when merging multi-source metadata, with interesting emerging solutions about relationship privacy preservation in social networking [43].

In addition to interoperability, when designing a horizontal MSN middleware, a critical issue to take into account is how to flexibly achieve a suitable tradeoff in terms of effectiveness and efficiency, depending on both the considered application domain and deployment environment. Awareness of resource state at different layers and abstraction levels, already used to optimize network-layer protocols [44], is clearly emerging as a central guideline to improve also MSN performance: for example, the performance of community detection and social multicast at the MSN middleware upper layer can be jointly optimized with the communication and routing functionalities of the networking layer. To mention another example, performance improvements based on lower-layer state visibility can be achieved in social metadata inference: a general MSN middleware should enable the required tradeoff between exhaustiveness of social graph inference (completeness, accuracy, precision, ...) and execution time depending on application-specific requirements. Indeed, completeness of inference results may depend not only on the adopted inference algorithms, but also on the spatial/temporal scopes and costs of acquisition/transfer of context to the node where the inference processing is run. Therefore, the tradeoff between the quality of inference results and performance could be improved by jointly operating on both the social inference and the context management facilities. For instance, less intensive inference algorithms could be adopted, such as approximate reasoning (anytime algorithms, approximate entailment) [45], and/or the scope of context acquisition could be extended/restricted depending on the dynamically required tradeoff.

To effectively support different MSN applications, another important technical challenge is the flexible definition of MSN middleware component allocation (where MSN middleware components execute and with which other components they interwork), typically only at deployment time. A variety of allocation schemas could be adopted, from completely centralized to hybrid and completely distributed ones. For instance, in a very simple centralized schema, the MSN middleware executes its facilities, e.g., for social data inference, community detection, and view management, on a single central server, typically owned and managed by the MSN application provider. As discussed for MobiSoc in Section 3, such a centralized management points has the advantage of facilitating complex social tie inference and security management. Of course, centralized allocation of middleware facilities may experience relevant performance penalties in the case of large numbers of concurrent users, large amounts of context metadata and application data to be moved, and unstable connection links between mobile users and the centralized server. In a completely distributed MSN solution, all middleware facilities run locally at users' mobile devices: the main weak point is the need of inter-node coordination to enable complex social tie inferences based on global knowledge of social graphs. In many deployment scenarios of practical interest, the MSN middleware deployment can follow hybrid allocation schemas where only some facilities, possibly working only on partial social views or with limited scopes, could be centralized. For instance, the architectural choice of whether to physically separate context metadata storing/management from reasoning is a key choice to take into account since the first phases of MSN middleware/application design, by carefully considering the different tradeoffs possible between achievable exhaustiveness and performance (e.g., communication-related energy consumption vs. data duplication).

A very crucial management challenge, which is worthwhile for a separated category per se, is scalability. As a general consideration, scalability is recognized as a central issue in all pervasive computing applications, including urban-wide smart city scenarios [46]. In MSN, there is the need to dynamically examine and consider huge amounts of continuously evolving context metadata and user-generated data to be shared, from an enormous set of possible sources. It is necessary to dig/infer/aggregate metadata results in limited time intervals, e.g., promptly exploit them while some users are co-located. These scalability issues, strictly interconnected with performance tradeoffs and allocation considerations, call for novel algorithms, heuristics, mechanisms, and tools capable of prioritizing the metadata to be considered and the collaboration opportunities possibly available. Often a very effective way to reduce this problem complexity is to make the MSN middleware/application focus only on the entities in the current locality of the focal node, with partial scopes, at least as a first and most common approximation (when feasible given the application domain and circumstances). Another challenging aspect of scalability is the dissemination of contents/metadata among the dynamic participants of a MSN view in an optimized way, by considering sparse connectivity, probability of opportunistic encounters, and limitations in mobile node resources, e.g., via probabilistic and socially-aware replication techniques [8,47].

As a final remark, it is worth pointing out that some of the aforementioned design considerations and technical issues are similar to the ones raised for the effective development and deployment of context-aware middleware solutions in general. This is not surprising given that social metadata can be seen as a specific but very relevant case of context metadata, associated with both "raw" data of explicit social relationships (such as the ones declared in Facebook social graphs) and ''advanced'' information inferred from recognized activities, history of interactions, and complex correlations. Some research activities have achieved relevant results on how to effectively design and implement context management middleware, especially in the sub-areas of context modeling, processing, reasoning, and aggregation, as well as of security and privacy [34–37]. However, we claim that MSN scenarios have unique characteristics that require specific and novel middleware solutions, especially for modeling/representing/inferring social ties and for creating effective mobile communities depending on focal nodes of interest, spatial/temporal scopes, and local MSN views. For instance, the need of limiting the horizon of the currently discoverable entities, for a given user in a specific application, is much stronger for MSN because of the relatively large set of entities with potential relationships (social relationship graphs may be very large) and the relatively limited time available to benefit from an opportunistic encounter. Therefore, existing context-aware middleware can offer an interesting and useful basis of effective mechanisms, algorithms, and protocols, but additional MSN-specific solutions need to be efficiently designed, implemented, and validated to achieve the goals of interoperability, multiple-layer optimization, effective allocation, and scalability.

5. Conclusive remarks and open challenges

This survey article has tried to provide a complete and fresh overview of the primary MSN middleware solutions in the literature. After providing an original taxonomy of state-of-the-art MSN middleware, based on the abstractions of focal nodes and scopes, the paper has aimed to extract and describe some primary lessons learned, design guidelines, and possible tradeoffs. These distilled guidelines should be useful to orientate the community of researchers and practitioners in the field when exploiting existing MSN middleware or designing new supports on top of existing mechanisms and facilities. Our general claim is that there is an evident need toward general-purpose middleware for MSN, to be flexibly specialized and set depending on specific MSN application requirements and characteristics of the targeted deployment environment.

Even if relevant results in the MSN field have already been achieved, with some open technical challenges that still call for comprehensive, effective, and integrated solutions. Among the others, we claim that security challenges, expansion to ''less-traditional'' application domains, and synergy with ''big data'' processing solutions represent three main open issues deserving attention in the near future MSN research.

First, MSN-specific security issues call for novel solutions. MSN applications manage context metadata that are very security-sensitive, from mobility traces and user preferences/activities to human relationships; and can be further used to infer even more sensitive users' context. In addition, MSN middleware/application providers must also respect the privacy of community members by preventing the disclosure of sensitive information [43]. Enforcing privacy and access control is further complicated in MSN with local spatial scopes, where context metadata and contents come from multiple sources and may be linked to other data according to statically unpredictable patterns. Traditional Enterprise and Digital Right Management mechanisms do not fit well these settings, where the enforcement of usage control policies cannot rely on the availability of centralized repositories that store and evaluate the policies to enforce. Nevertheless, appropriate privacy/access/usage control policies should be defined and enforced over social data exchange: for instance, users should be able to easily specify both the data to protect and the set of peers allowed to access that data in a socially-aware manner (e.g., ''only let my colleagues know my location during weekdays from 9 am to 5 pm''). Existing work on social-aware policies for mobile environments [48,49] and policies to control dataflows in loosely connected networks [50] offers a promising approach. In addition, it is worth noticing that there are also open questions related to data ownership [51,52]. Who owns and control how social metadata available on MSN middleware can be exploited and distributed?[4] This technical challenge is perceived as a critical aspect to address in MSN platforms, especially when centralized context management and inference are adopted. Decentralized MSN solutions could be preferable under the perspective of protecting data ownership: in this case no single entity can have easy access to all context data (albeit encrypted) to infer relationships or behaviors. That motivates the recent push (e.g., from institutions, regulators, and users' communities) toward decentralized MSN middleware, whose success may change also the business model behind the provisioning of OSN applications, forcing providers to abandon the current model of value/gain connected to the number of users and the full ownership of their context metadata.

Second, the expansion of MSN solutions to many application domains, even far from what was expected at the beginning, is starting to pose more critical and demanding requirements on MSN middleware and the offered quality. MSN application areas already include, for instance, healthcare assistance and emergency response services, where context metadata are dynamically exploited to build adequate, ad-hoc, and impromptu communities to help people by properly integrating the needed skills and attitudes. These different application domains will make differentiated and application-specific facilities emerge, to be properly built on top of the previously described general-purpose MSN middleware. In addition, they will result in challenging and relevant testbeds to validate the general applicability of MSN middleware (and of its flexibility in terms of different tradeoffs achievable).

Third, there is a very interesting and cross-disciplinary trend toward the collection of enormous amounts of social metadata to build more appropriate and accurate models of social behaviors/trends and their evolutions. For instance,

[4] This issue has been already considered a privacy breach recently when users discovered that their profile data are still kept on some OSN applications even after their account removal.

consider the ambitious objectives of the EU FuturICT knowledge accelerator coordination action [53], which aims at considering world-wide social-life observations to understand social behaviors, to determine rising trends (e.g., in the Living Earth Simulator), and to predict complex events with sufficient advance time (e.g., in the Economic Crisis Observatory). On the one hand, MSN middleware could represent a valuable and irreplaceable source of (partially pre-processed) social-life metadata. On the other hand, these ambitious application scenarios exacerbate the related technical challenges, especially scalability, pushing for (i) the organization of MSN overlays in dynamic hierarchical localities, to aggregate and reduce inter-locality metadata transfer; (ii) lightweight approximated techniques (possibly social-aware in their turn) for extracting probabilistic metadata trends; and (iii) the exploitation of external resources (e.g., cloud computing) where to properly offload heavy processing operations, especially off-line determination of complex social ties and patterns.

References

[1] N.D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, A. Campbell, A survey of mobile phone sensing, IEEE Communications Magazine 48 (9) (2010) 140–150.
[2] M. Conti, S. Giordano, M. May, A. Passerella, From opportunistic networks to opportunistic computing, IEEE Communications Magazine 48 (9) (2010) 126–139.
[3] M. Conti, S.K. Das, C. Bisdikian, M. Kumar, L.M. Ni, A. Passarella, G. Roussos, G. Tröster, G. Tsudik, F. Zambonelli, Looking ahead in pervasive computing: challenges and opportunities in the era of cyber-physical convergence, Pervasive and Mobile Computing 8 (1) (2012).
[4] Z. Zhang, Feeling the sense of community in social networking usage, IEEE Transactions on Engineering Management 57 (2) (2010) 225–239.
[5] D. Katsaros, N. Dimokas, L. Tassiulas, Social network analysis concepts in the design of wireless Ad Hoc network protocol, IEEE Network 24 (6) (2010) 23–29.
[6] C. Boldrini, M. Conti, F. Delmastro, A. Passarella, Context- and social-aware middleware for opportunistic networks, Journal of Network and Computer Applications 33 (5) (2010) 525–541. Elsevier.
[7] C. Boldrini, M. Conti, A. Passarella, Design and performance evaluation of content place, a social-aware data dissemination system for opportunistic networks, Computer Networks 54 (4) (2010) 589–604. Elsevier.
[8] N. Kayastha, D. Niyato, P. Wang, E. Hossain, Applications, architectures, and protocol design issues for mobile social networks: a survey, Proceedings of the IEEE 99 (12) (2011) 2130–2158.
[9] M. Conti, M. Kumar, Opportunities in opportunistic computing, IEEE Computer 43 (1) (2010) 42–50.
[10] Q. Jones, S.A. Grandhi, P3 systems: putting the place back into social networks, IEEE Internet Computing 9 (5) (2005) 38–46.
[11] E.F. Churchill, C.A. Halverson, Social networks and social networking, IEEE Internet Computing 9 (5) (2005) 14–19. Guest editors' introduction.
[12] C. El Morr, J. Kawash, Mobile virtual communities research: a synthesis of current trends and a look at future perspectives, Inderscience International Journal of Web Based Communities 3 (4) (2007) 386–403.
[13] E.M. Daly, M. Haahr, Social network analysis for routing in disconnected delay-tolerant manets, in: ACM Int. Symp. Mobile Ad Hoc Netw. Comput., 2007, pp. 32–40.
[14] E.M. Daly, M. Haahr, Social network analysis for information flow in disconnected delay-tolerant MANETs, IEEE Transactions on Mobile Computing 8 (5) (2009) 606–621.
[15] N. Eagle, A. Pentland, D. Lazer, Inferring social network structure using mobile phone data, Proceedings of the National Academy of Sciences, PNAS 106 (36) (2007) 15274–15278.
[16] A.T. Campbell, S.B. Eisenman, N.D. Lane, E. Miluzzo, R.A. Peterson, H. Lu, K. Zheng, M. Musolesi, K. Fodor, A. Gahng-Seop, The rise of people-centric sensing, IEEE Internet Computing 12 (4) (2008) 12–21.
[17] B.J.F. van Beijnum, P. Pawar, C.B. Dulawan, H.J. Hermens, Mobile virtual communities for telemedicine: research challenges and opportunities, International Journal of Computer Science and Applications, Technomathematics Research Foundation 6 (2) (2009) 19–37.
[18] J. Subercaze, P. Maret, J. Calmet, P. Pawar, A service oriented framework for mobile business virtual communities, in: FIP TC 5 WG 5.5 9th Working Conf. Virtual Enterprises, Vol. 283, Springer Verlag, Poznan, Poland, 2008, pp. 493–500.
[19] M. Terry, E.D. Mynatt, K. Ryall, D. Leigh, Social net: using patterns of physical proximity over time to infer shared interests, in: Int. Conf. on Human Factors in Computing Systems, ACM, Minneapolis, 2002, pp. 816–817.
[20] N. Eagle, A. Pentland, Social serendipity: mobilizing social software, IEEE Pervasive Computing 4 (2) (2005) 28–34.
[21] S. Boehm, J. Koolwaaij, M. Luther, B. Souville, M. Wagner, M. Wibbels, Introducing IYOUIT, in: 7th Int. Semantic Web Conf., in: LNCS, vol. 5318, Karlsruhe, 2008, pp. 804–817.
[22] J. Haikio, T. Tuikka, E. Siira, V. Tormanen, Would you be my friend? Creating a mobile friend network with 'hot in the city, in: IEEE 43rd Hawaii Int. Conf. on System Sciences (HICSS), IEEE Computer Society, 2010, pp. 1–10.
[23] A. Beach, M. Gartrell, S. Akkala, J. Elston, J. Kelley, K. Nishimoto, B. Ray, S. Razgulin, K. Sundaresan, B. Surendar, M. Terada, R. Han, WhozThat? Evolving an ecosystem for context-aware mobile social networks, IEEE Network 22 (4) (2008) 50–55.
[24] M. Tran, J. Han, A. Colman, Social context: supporting interaction awareness in ubiquitous environments, in: 6th Annual Int. Conf. on Mobile and Ubiquitous Systems: Networking & Services, IEEE Computer Society, Toronto, 2009, pp. 1–10.
[25] A. Gupta, A. Kalra, D. Boston, C. Borcea, MobiSoC: a middleware for mobile social computing applications, Mobile Networks and Applications Journal 14 (1) (2009) 35–52. Kluwer Academic Publishers.
[26] D. Bottazzi, R. Montanari, A. Toninelli, Context-aware middleware for anytime, anywhere social networks, IEEE Intelligent Systems 22 (5) (2007) 23–32.
[27] A. Pietilainen, E. Oliver, J. LeBrun, G. Varghese, C. Diot, MobiClique: middleware for mobile social networking, in: 2nd ACM Workshop on Online Social Networks, New York, 2009, pp. 49–54.
[28] A. Toninelli, A. Pathak, V. Issarny, Yarta: a middleware for managing mobile social ecosystems, in: 6th Int. Conf. on Advances in Grid and Pervasive Computing, in: LNCS, vol. 6646, Oulu, 2011, pp. 209–220.
[29] V. Arnaboldi, M. Conti, F. Delmastro, Implementation of CAMEO: a context-aware middleware for opportunistic mobile social networks, in: IEEE Int. Symp. World of Wireless, Mobile and Multimedia Networks, IEEE Computer Society, 2011, pp. 1–3.
[30] T. Nguyen, S.W. Loke, T. Torabi, H. Lu, PlaceComm: a framework for context-aware applications in place-based virtual communities, IOS Journal of Ambient Intelligence and Smart Environments 3 (1) (2011) 51–64.
[31] D. Brooker, T. Carey, I. Warren, Middleware for social networking on mobile devices, in: 21st IEEE Australian Software Engineering Conf., IEEE Computer Society, 2010, pp. 202–211.
[32] Friend of a friend. URL: http://www.foaf-project.org.
[33] OpenSocial. URL: http://www.opensocial.org/page/opensocial-foundation-faq.
[34] K. Henricksen, J. Indulska, Developing context-aware pervasive computing applications: models and approach, Pervasive and Mobile Computing 2 (2) (2006) 37–64.
[35] M. Baldauf, S. Dustdar, F. Rosenberg, A survey on context-aware systems, Inderscience International Journal of Ad Hoc and Ubiquitous Computing 2 (4) (2007) 263–277.

[36] C. Bettini, O. Brdiczka, K. Henricksen, J. Indulska, D. Nicklas, A. Ranganathan, D. Riboni, A survey of context modelling and reasoning techniques, Elsevier Pervasive and Mobile Computing Journal 6 (2) (2010) 161–180.
[37] N. Roy, G. Tao, S.K. Das, Supporting pervasive computing applications with active context fusion and semantic context delivery, Elsevier Pervasive and Mobile Computing Journal 6 (1) (2010) 21–42.
[38] O. Yoosoo, H. Jonghyun, W. Woontack, A context management architecture for large-scale smart environments, IEEE Communications Magazine 48 (3) (2010) 118–126.
[39] M. Roman, C. Hess, R. Cerqueira, R.H. Campbell, K. Nahrstedt, Gaia: a middleware infrastructure to enable active spaces, IEEE Pervasive Computing Magazine 4 (1) (2002) 74–83.
[40] O. Riva, Contory, a middleware for the provisioning of context information on smart phones, in: ACM/IFIP/USENIX International Conference on Middleware, in: LNCS, vol. 4290, Springer Verlag, 2006, pp. 219–239.
[41] J. Veijalainen, Developing mobile ontologies; who, why, where, and how? in: 8th Int. Conf. on Mobile Data Management, IEEE Computer Society, Mannheim, 2007, pp. 398–401.
[42] Artemis SOFIA project. URL: http://www.sofia-project.eu/.
[43] N. Li, N. Zhang, S.K. Das, Preserving relation privacy in online social network data, IEEE Internet Computing 15 (3) (2011) 35–42.
[44] V. Srivastava, M. Motani, A cross-layer design: a survey and the road ahead, IEEE Communications Magazine 43 (12) (2005) 112–119.
[45] S. Rudolph, T. Tserendorj, P. Hitzler, What is approximate reasoning? in: 2nd Int. Conf. on Web Reasoning and Rule Systems, in: LNCS, vol. 5341, Karlsruhe, 2008, pp. 150–164.
[46] NSF Workshop on Pervasive Computing at Scale, Seattle, 2011. URL: http://sensorlab.cs.dartmouth.edu/NSFPervasiveComputingAtScale/.
[47] S. Ioannidis, A. Chaintreau, L. Massoulie, Distributing content updates over a mobile social network, Mobile Computing and Communications Review 13 (1) (2009) 44–47.
[48] A. Toninelli, O. Lassila, D. Khushraj, R. Montanari, What's on users' minds? Toward a usable smart phone security model, IEEE Pervasive Computing 8 (2) (2009) 32–39.
[49] S. Hachem, A. Toninelli, A. Pathak, V. Issarny, Policy-based access control in mobile social ecosystems, in: IEEE Int. Symp. on Policies for Distributed Systems and Networks, IEEE Computer Society, Pisa, 2011, pp. 57–64.
[50] G. Russello, E. Scalavino, N. Dulay, E.C. Lupu, Coordinating data usage control in loosely-connected networks, in: IEEE Int. Symp. on Policies for Distributed Systems and Networks, IEEE Computer Society, Washington, 2010, pp. 30–39.
[51] D. Rosenblum, What anyone can know: the privacy risks of social networking sites, IEEE Security and Privacy 5 (3) (2007) 40–49.
[52] S. Buchegger, D. Schiöberg, L. Vu, A. Datta, PeerSoN: P2P social networking: early experiences and insights, in: ACM Second EuroSys Workshop on Social Network Systems, New York, Proceeding ACM, 2009, pp. 46–52.
[53] Future ICT coordination action. URL: http://www.futurict.ethz.ch/FuturICT.