



Contents lists available at SciVerse ScienceDirect

Pervasive and Mobile Computing

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

Fast track article

Looking ahead in pervasive computing: Challenges and opportunities in the era of cyber–physical convergence

Marco Conti^{a,*}, Sajal K. Das^b, Chatschik Bisdikian^c, Mohan Kumar^b, Lionel M. Ni^d,
 Andrea Passarella^a, George Roussos^e, Gerhard Tröster^f, Gene Tsudik^g, Franco Zambonelli^h

^a IIT-CNR, Pisa, Italy^b University of Texas at Arlington, USA^c IBM Research, USA^d HKUST, Hong Kong^e Birkbeck College, University of London, UK^f ETH Zurich, Switzerland^g University of California at Irvine, USA^h University of Modena e Reggio, Italy

ARTICLE INFO

Article history:

Received 10 May 2011

Received in revised form 23 September 2011

Available online 28 October 2011

Keywords:

Pervasive computing

Cyber–physical convergence

Self-*

Social networks

Wearable computing

Opportunistic networking and computing

Data storage

Quality of Information

Cyber–world security

ABSTRACT

The physical environment is becoming more and more saturated with computing and communication entities that interact among themselves, as well as with users: virtually everything will be enabled to source information and respond to appropriate stimuli. In this technology-rich scenario, real-world components interact with cyberspace via sensing, computing and communication elements, thus driving towards what is called the *Cyber–Physical World* (CPW) convergence. Information flows from the physical to the cyber world, and vice-versa, adapting the converged world to human behavior and social dynamics. Indeed humans are at the center of this converged world since information about the context in which they operate is the key element to adapt the CPW applications and services. Alongside, a new wave of (human) social networks and structures are emerging as important drivers for the development of novel communication and computing paradigms. In this article we present some of the research issues, challenges and opportunities in the convergence between the cyber and physical worlds. This article is not a comprehensive survey of all aspects of the CPW convergence. Instead, it presents some exciting research challenges and opportunities identified by members of the journal's editorial board with a goal to stimulate new research activities in the emerging areas of CPW convergence.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

As originally envisioned by Mark Weiser about 20 years ago, pervasive devices and services abound permeating and becoming parts of our daily lives [1–3]. Smartphones, e-readers, GPS-enabled cameras, tablet computers, and other gadgets are already having a transformative effect on the development of our ecosystems by interlinking the cyber and physical worlds. By exploiting these devices and various technologies, information about physical reality (e.g., collected through sensor nodes) is seamlessly transferred into the cyber world where it is elaborated to adapt cyber applications and services to the physical context, and thus possibly modifying/adapting the physical world itself through actuators, see Fig. 1. This

* Corresponding author.

E-mail address: marco.conti@iit.cnr.it (M. Conti).

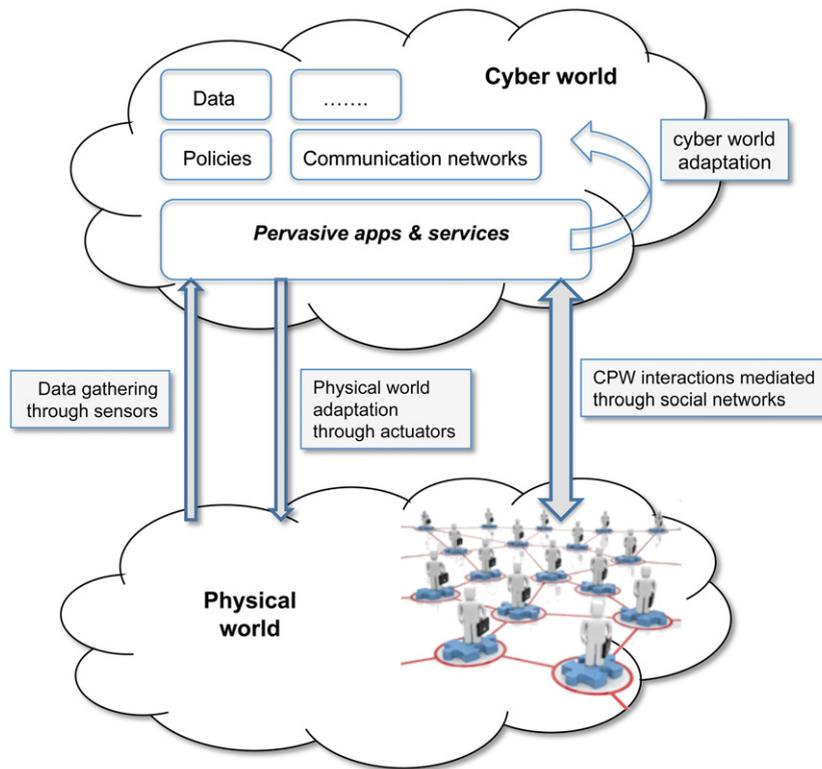


Fig. 1. Physical and cyber world interactions.

opens up the space for the creation of innovative services to better understand and interact with the surrounding physical world as well as with the social activities therein, and more generally to promote capabilities of “augmented sensing and interacting” within the physical and social worlds [4,5]. The vision of a virtual world overlaid on the physical world to continuously monitor it and possibly take intelligent actions to adapt the cyber world (applications and devices) to our needs is part of our common view of a pervasive or ubiquitous computing system. However, the *Cyber-Physical World* (CPW) convergence is opening up new research directions for pervasive computing researchers. Indeed, in a converged world, actions and information produced in the physical world can affect and modify the personal and social contexts, which may then affect how information and services are handled in the cyber world. This latter aspect, which needs more in-depth investigations, is emerging as an extremely challenging research area. Clearly, humans are at the core of the CPW convergence; each person has access to several devices of various connectivity and computing capabilities through which he/she can interact with the virtual world, thus linking the physical world and the electronic world of users’ devices. By contributing to set up the cyber-world infrastructure with the help of data, communications and devices, the users (and hence the physical world) can also have a major impact on the organization of the cyber infrastructure.

Human social structures can play a central role in controlling the way information spreads in the cyber world and thus how pervasive applications and services can be created and orchestrated. As shown in Fig. 2, by translating human relationships into the cyber world, we embed in the electronic devices the social relationships that enable humans to effectively handle and share large amounts of information. The resulting network is named *electronic social network*. Human social networks exhibit remarkable dynamism and structural properties that may significantly affect the quality of the information (e.g., trust and reputation, relevance, reliability, etc.), and also the way information may circulate both in the physical and cyber worlds.

On the other hand, by extending human social networks with cyber-world social networks,¹ we possibly modify their structure/organization, and in turn affect the way humans share information in the physical world. Furthermore, since information is increasingly generated by users and/or is available on users’ devices, human social networks play a very important role in accessing and circulating the massive quantities of information present in the cyber world by establishing overlays that link together users and their surrounding environment (i.e., components in the physical world) with users’ devices (i.e., components in the cyber world). So far, this aspect has received relatively less attention in the research community and hence been not well investigated although its importance is significantly increasing at an ever faster rate.

¹ Cyber-world social networks include both the electronic social networks of human devices and online social networks like Facebook.

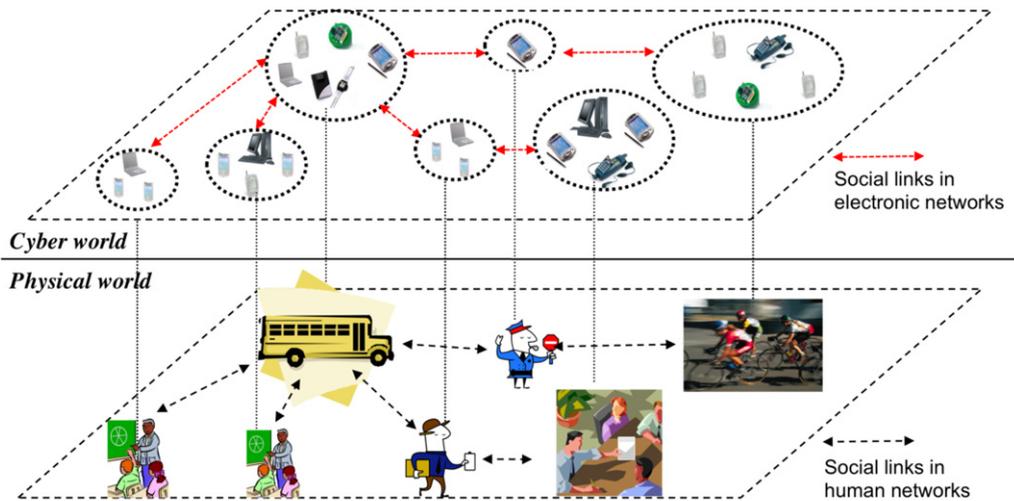


Fig. 2. Human and electronic social relationships.

In this paper we attempt to analyze the emerging research issues, challenges and opportunities in the pervasive and mobile computing domain, arising due to the tight interactions between the physical and cyber worlds.

In the convergent CPW, a wide variety of smart devices, such as RFID tags, sensors and actuators, sensor-rich smartphones, or proximity sensing technologies, are leading to the emergence of an integrated and very dense infrastructure for monitoring the physical world, and hence collecting information related to user behaviors, their requirements and dynamics. The management of such a complex infrastructure calls for the development of effective and scalable policies for handling the cooperation among these devices and adapting their behavior to the rapidly changing physical (e.g., location and other environmental characteristics) and social (e.g., current activity and mutual interaction) contexts from which the digital services (e.g., information access and personal communications) are invoked. Developing architectures and policies that are able to autonomously adapt the components in the cyber-world is a major research area, as discussed in Section 2.

As mentioned earlier, the availability of a virtual infrastructure overlaid on the physical world can provide continuous interactions between the two worlds. To this end, Section 3 contends that the cyber world infrastructures can be exploited for studying the physical world with a special attention to achieving a better understanding of human behavior to adapt pervasive computing applications and services to the users' needs. Additionally, the virtual world infrastructure can also modify the physical world by exploiting actuators/actors distributed in the latter. For example, the information collected from the physical world can be used to provide ubiquitous services tailored to the user context and needs, such as GPS navigator systems providing real-time traffic information and updating routes accordingly, or smart phones informing us about the current locations of our friends.

Section 4 describes how human social relationships will be widely exploited (at different levels with different scales) in the virtual world in order to develop new computing and communication paradigms. Specifically, as investigated in the Socialnets project,² *physical social relations* can be defined between pervasive computing devices, by exploiting the social links existing among their users. By following the physical interaction among the users, the information circulates among their devices through one-hop wireless connections. In fact, driven by their social links, the users (and hence their devices) often come in close proximity with the members of their social networks. This is the basis for a new communication paradigm called *opportunistic networking* in the sense that the networking protocols opportunistically exploit the social links existing among devices for data exchange [6–8]. Opportunistic networking thus represents a step towards *mobile social networking*. In addition to physical relations among users, the human social structures also have a *virtual dimension*, in which a social network is defined by the common relevance of information to users. The participants to this social structure are not necessarily physically co-located but are tied by common interests or human behavior. Again such virtual social links are a powerful tool for the dissemination of information. For example, online social networks utilize virtual social structures that link together people sharing common interests and/or members of the same social group(s). A further step to this view is exploiting the “social” cooperation among devices for building service overlays on top of mobile social networks. These overlays support sharing and composition of the computing services available in the users' devices, which can thus exploit the services available in cyberspace by offering much richer services to their users with respect to what is available on any individual device in isolation. This novel computing paradigm is currently referred to as *opportunistic computing* [7,9], and is discussed in Section 4.2.

² <http://www.social-nets.eu/>.

A key element that characterizes the CPW convergence is the huge volume of data, coming from many heterogeneous sources, that is flowing from the physical to the cyber world (and vice versa) and that requires appropriate handling. Data needs to be processed (e.g., data cleaning and calibration), properly organized (e.g., novel data models have to be identified) and stored. Furthermore, the quality of the data coming from these heterogeneous sources has to be evaluated with respect to the information requirements of the applications. To this end, the *Quality of Information* (QoI), i.e., the ability to judge available information for a given purpose, will play an important role. To be precise, QoI represents the basis for searching and selecting the most appropriate information source(s) against the information needs of the applications. These aspects are discussed in Section 5. Specifically, Section 5.1 analyzes the challenges associated with data storing, while Section 5.2 discusses the quality of the information and the associated contexts.

The fact that a large number of small, inexpensive, and resource-constrained computing devices are pervasively embedded in the physical world implies that novel security and privacy solutions must be devised [10]. This is particularly important when dealing with human social contexts and vital information [11]. Section 6 discusses some of the security challenges for devices operating in the cyber-world.

Finally, Section 7 concludes the paper and also discusses additional research challenges.

2. Autonomic behavior of the cyber world infrastructure³

The CPW convergence scenario is characterized by a large number of mobile devices that the users bring around or that are spread into the environment. The components of the cyber infrastructure should be able to dynamically and opportunistically connect and interact with each other, and to adapt themselves to the rapidly changing physical/social world.

To guarantee the cooperation among these devices and their ability to adapt to changing scenarios, a well recognized research challenge relates to the need for enforcing autonomic, self-managing and self-adaptive behavior, both at the level of the infrastructure and that of its services [12]. Indeed, the CPW convergence scenario is characterized by increasing decentralization and dynamics, including a variety of highly distributed devices of an ephemeral and/or mobile nature, with users and developers capable of injecting new components and services at any time. This can make it impossible for developers and system managers to stay in the control loop and directly intervene in the system for configuration and maintenance activities (e.g., for configuring or personalizing new distributed devices or services, for fixing problems in existing devices and services, or for optimizing resource exploitation).

Accordingly, in recent years, significant efforts have been made at both infrastructure and service levels, to promote autonomic and adaptive behavior in pervasive computing systems. However, in our opinion, most of the current proposals suffer from several limitations when dived in future scenarios. For this reason, in Section 2.1 we analyze and discuss the major limitations in the current approaches, while in Section 2.2 we present the emerging research challenges to tackle these limitations.

2.1. Limitations of current approaches to autonomic adaptation

A number of autonomic computing approaches propose “add-ons” to be integrated in existing frameworks such as the one from “à la IBM”, that suggests coupling sophisticated control loops to existing systems to support self-management [13]. The result is often an increased complexity of current frameworks, which definitely does not suit the lightweight characteristics of CPW infrastructures, a requirement due to the limited capabilities of pervasive and mobile devices and inherent decentralization of the CPW scenarios.

A number of other proposals indeed suggest relying on light-weight and fully decentralized approaches, typically inspired by natural phenomena of self-organization and self-adaptation [14,12]. However, despite the promises and potentials of nature-inspired approaches, most of the existing proposals exploit the natural inspiration only for the implementation of specific algorithmic solutions or for realizing specific distributed services, at the infrastructure or the user level, rather than tackling the issue of autonomic self-adaptation in a comprehensive way.

In addition, although some recent research efforts have focused on the social aspects of pervasive computing and on the provisioning of innovative social services, they do not properly account for the social level as an indistinguishable dimension of the overall pervasive computing fabric. That is, they do not account for the fact that users, other than simply consumers or producers of services, are integral components of the overall infrastructure, and contribute to it via human sensing, actuating, and computing capabilities [15]. Thus, “in isolation” approaches to autonomic adaptation cannot simply fit the emerging CPW scenarios or their convergence.

2.2. Emerging challenges

Based on the above considerations, our belief is that – rather than looking for one-off solutions to specific adaptation problems from specific limited viewpoints – there is a need to deeply rethink the modeling and architecting of future

³ By Franco Zambonelli.

pervasive computing systems. As challenging as this can be, one should try to *account for a foundational and holistic way to tackle the complex needs of adaptation and autonomic behavior* of such systems. The final goal should be that of making such systems inherently capable of autonomic self-management and adaptation at the collective level, and having the distinction between infrastructure, services, and social levels blur or vanish completely.

The accomplishment of this broad goal poses the following research issues and challenges from a global perspective.

- *Comprehensive situation-awareness.* The need for context-awareness is a recognized issue in pervasive computing over the years. However, in recent years, the specific challenges associated with context-awareness are notably changing. On the one hand, acquiring the data necessary to support context-awareness is no longer an issue. In fact, as analyzed in Section 5, the opposite issue of handling an overwhelming amount of data is arising. Future scenarios will require autonomous adaptation activities to be driven by more comprehensive levels of awareness than traditionally enforced in context-aware computing models [16]. Most current approaches to context-awareness leave up to each component of the system to access and digest the information required for adaptation decisions, but this task can become overwhelming if awareness has to encompass both situations occurring at the many different levels of the system and the strict locality of components. Thus, a challenging research issue is defining novel tools to provide components of the pervasive computing infrastructure with expressive and compact representations of complex multi-faceted situations, so as to effectively drive each and every activity of the components in a collectively coordinated way [17].
- *Top-down vs. Bottom-up.* Along with the traditional “top-down” approaches to the engineering of pervasive computing systems, in which specific functionalities or behavior are achieved by explicit design, “bottom-up” approaches are also being proposed and adopted to achieve functionalities via spontaneous self-organization, as it happens in natural systems. Most likely, both approaches will co-exist in future pervasive computing scenarios, the former being applied to the engineering of specific local functionalities, while the latter being applied to the engineering of large-scale behaviors (or simply emerging as natural collective phenomena of the system). Yet, understanding the trade-offs between the power of the top-down and bottom-up adaptations, and studying how the two approaches can co-exist and possibly conflict in future systems, is an open research challenge. Facing such a challenge is necessary to smooth the tension between the two approaches and their effective applications [18].
- *The power of the masses.* As pervasive computing infrastructures are becoming very large scale sources of huge amounts of data, involving a large number of users, it becomes necessary to understand the “power of the masses” principle as far as the participatory pervasive computing processes of collective adaptations are involved. Most phenomena and mechanisms of self-adaptation and self-organization are now typically studied in small-scale systems. Future CPW convergence scenarios will involve billions of users, devices, and data items. This will make it necessary to understand how and to what extent – at such large scales – such phenomena can express forms of adaptation and situation-awareness (or, observable “intelligence”) much more effectively than those obtained today with more traditional forms of distributed computing and artificial intelligence techniques [19].
- *Decentralized control.* Strictly related to the above, the existence and exploitation of phenomena of collective adaptation must necessarily come along with models and tools, thus making it possible to control “by design” the overall behavior of the pervasive systems and its sub-parts [20]. Clearly, due to the inherent decentralization of the systems, such control tools should be decentralized too. Thus, the challenging issue of laying out foundations and defining tools for the decentralized control of complex CPW scenarios will arise. In addition, to understand if such control is effective, it will be necessary to identify suitable “measures” to characterize the behavior of the systems under control. The issue of defining sound measures for future CPW scenarios can itself be a challenging research area, given the size of the target scenarios and the fact that they will exist not to serve a single well-defined (and thus easily measurable) objective, but to serve the many diverse requirements of many actors.
- *Diversity and evolvability.* As of now, most studies related to the collective adaptation focus on systems with a limited and fixed number of component classes. However, this is far from even approximating the increasing diversity of components (devices and services) of future and emerging CPW convergence scenarios, and their continuous evolution. Promoting diversity and evolvability is essential for user contributions and personalization, and hence to support a continuous process of value co-creation in the overall pervasive infrastructures. Yet, besides being potential sources of complexity, diversity and evolvability are also recognized to play a key role in collective adaptation and intelligence in complex ecosystems and organisms, respectively. Thus, it will be necessary to understand how these issues can be accommodated and controlled, possibly by shifting the current focus from biological forms of collective adaptation to ecological and social ones.
- *Mechanisms design.* To get the most from future CPW scenarios, and provided that the technology for autonomic and opportunistic networking and computing will be there (see Section 1), there must be means to ensure that all available components (devices, services, human actors) are prone to put their capabilities for the collective service to opportunistically and effectively interact with each other. This calls for identifying and designing novel interaction mechanisms to incentivize and support adaptive interactions among the many and diverse components of the CPW system. Several approaches, typically based on auction mechanisms, have already been proposed to promote cooperative and/or opportunistic sharing of sensing devices [21,22]. Most likely, though, future CPW convergence scenarios will require identifying brand new incentive and pricing mechanisms.

In summary, the need for laying out novel foundations for the modeling and architecting of autonomic and self-adaptive CPW systems and infrastructures opens up a large number of important research challenges. The short list provided

above is far from being exhaustive. In particular, the list does not account for the many inter-disciplinary issues that the understanding of large-scale socio-technical organisms (as future CPW systems will be) and their collective adaptive behaviors involve [23], thereby exploiting the modern lessons of applied psychology, sociology, social anthropology, and macro-economy, in addition to those of systems biology, ecology and complexity science.

3. Studying the physical world from the cyber world

The knowledge of the human behavior is a fundamental step for tuning the pervasive applications to the users' context. Furthermore, a better understanding of the human behavior opens up new possibilities in several areas of science and engineering, e.g., urban planning, transportation, energy management, etc.

In Section 3.1 we discuss the research challenges in using pervasive computing methods and tools for studying the human behavior, while Section 3.2 focuses on the wearable computing paradigm, which is a fundamental tool for continuously monitoring human behavior.

3.1. Exploring human dynamics and social networks with pervasive computing techniques⁴

The transparent interlinking of cyber and physical reality, a core ingredient of pervasive computing, also enables the automatic observation and measurement of human behavior and in particular reveals the pattern of interactions between individuals at a scale and level of detail that is simply unprecedented in the history of humankind. The opportunities presented by this capability are numerous: Information captured in this way allows for the analysis, modeling and experimentation with human behaviors at the individual as well as social levels. It also supports data-driven dynamic adaptation of systems in response to inferences derived from such data and the observed user habits, preferences and routines as they emerge. As a consequence, research in pervasive computing is presented with a unique opportunity to understand and manage the dynamics of human behavior in order to make pervasive computing systems more usable and tractable.

The study of human dynamics through observation by pervasive computing systems has emerged as a common thread across several recent research projects. For example, the Dartmouth Campus project explored the patterns of students and staff movement on the university grounds using records of device associations to wireless network infrastructure [24]; Cityware explored the use of short-range wireless networking technology to track and cluster individual behaviors and identify distinct communities in the city [25]; Urban Tapestries explored the dynamics of participation within a pervasive user-generated media authoring system [26]; Senseable City visualizes collective human behavior captured through mobile network traffic records overlaid on metropolitan urban structures [27]; Hagggle and Socialnets use personal encounters as the main primitive for the development of novel models for networking [28,29]; and recent work at the *Center for Complex Network Research* employed complex network theory to analyze mobile phone location records to explore the structure of social networks [30]. Moreover, a recent special issue of this journal reported on a variety of investigations in the same spirit [31]. These and other projects have initiated *the study of methods and techniques that can be used to investigate the structure and evolution of the dynamics of human behavior under the lens of pervasive computing*. These early explorations inevitably focus on only a subset of the full gamut of possible life-long data, alternatively considering location, activities, proximity to others, interaction with artifacts, or environmental situations, thus offering only a partial view of the full range of possible dynamics. Such data sets can be significantly expanded with related behavioral attributes and more critically, they can be combined with the so-called Digital Footprint data, such as email exchanges, browsing, buying, authoring, social networking, sharing content and so on. The combination of the two allows for much more effective analyses, which can lead to more sophisticated and comprehensive representations of reality.

To be sure, the capability to conduct mass observation at this scale offers novel opportunities not only for computing but also for urban planning, transportation, energy management and several other areas of engineering and the sciences. For example, it is increasingly argued that the mass observation afforded by pervasive computing is a major factor behind the emergence of computational social science, which is redefining methods and techniques for the investigation of a wide variety of social phenomena [32]. For pervasive computing, the role of such interdisciplinary research is far greater than at any previous time during the evolution of information and communication technology because human behaviors are deeply embedded within its systems and indeed many aspects of their operation are dictated by the constraints set by human physicality. Indeed, placing humans as system-level participants within computing ecosystems rather than viewing them simply as users often requires a set of research skills that extends beyond the current reach of computing. To create rich representations of human dynamics extending beyond elementary descriptive statistical frameworks of observed behaviors, it is necessary to understand human agency in all its different expressions, an observation that calls for a cross-disciplinary programme of research.

One implication of the size of the captured data sets for future explorations into human dynamics is that they will likely necessitate the application of so-called Big Data approaches (see Section 5 for a more detailed discussion of data management in this context). The first implication of this observation is that it is necessary to develop suitable data

⁴ By George Roussos.

management techniques and algorithms so that analyses can fully benefit from the data sets available, which can quickly grow to the order of petabytes in full-scale systems. But more data also bring about qualitative differences to the possible methodological approaches. For example, the availability of very large data sets makes much more attractive the use of computers to automatically learn a model from the data rather than maintain the traditional approach which reserves the model-building role for a specialist [33]. The data-driven approach appears to be particularly well suited to complicated and unpredictable domains such as human behavior, which are unlikely to be fully described by a simple set of equations. Recent evidence in fact suggests that employing even relatively simple non-parametric density approximation models on such data sets can provide significant advantage [34]. Nevertheless, this methodology also has limitations in that such models are often complex and difficult to understand and do not necessarily match human intuition.

One issue that becomes immediately clear when discussing datasets of this type and size is that they set special challenges to the privacy and security of individuals as well as to trust in technology and personal identity. This is a complex problem that extends well beyond the specific algorithms or techniques used and attitudes towards it are still evolving as human societies are gaining a better understanding of the costs and benefits involved in the adoption of specific pervasive computing technologies. In some cases, even when data are carefully made anonymous, it is not uncommon for individuals to still be identifiable and targeted purely on behavioral information [35]. Finding the right balance between the desire to study human dynamics and the rights of the individual is a challenging task that has to be further debated.

Nevertheless, the majority of current work in this area is primarily exploratory in nature and has been restricted to one or a small subset of factors, for example either temporal or spatial behavior, which are rarely linked to semantic information. One approach that has been under-explored in addressing this limitation is the use of collective intelligence, which has been particularly successful in revealing and exploiting the dynamics of the web [36]. Recall that human dynamics fundamentally depends on aspects of human behavior that are internalized, such as beliefs, expectations, memories and agendas, and which traditional computing research rarely visits. Progress in this area would require cross-pollination of ideas with other scientific and social-scientific domains.

One direct application of human dynamics is in locating and navigating to the specific place, device, software component, person or object most relevant to a specific situation. This is a critical feature for effective operation of pervasive computing systems but also a task of considerable complexity. Recent work in this area has to be extended with further development of learning techniques for analysis and prediction [25]. In the longer-term, such work should lead to a comprehensive model providing for the spatial, temporal and semantic elements of particular situations. These methods would also be useful to support adaptation and are intimately related to the provision of context-awareness in applications. Capturing human dynamics can also reveal the characteristic patterns of mobility of users and associated devices and services, which can subsequently be employed to provide increased robustness and resilience of next generation networks [37]. Mobility due to human movement in particular can improve wireless network performance when seen under the so-called participatory sensing or opportunistic networking approaches [9]. In the long run, understanding typical patterns of behavior within specific situations would allow such behaviors to be viewed as predictable and thus be incorporated as a feature of network infrastructure [28,29]. Within this view of human activity, of particular significance are models of social behaviour [25]. For example, social network growth models can provide appropriate abstractions to understand the evolution and core properties of pervasive networks. Such an understanding can lead in the longer run towards the development of effective information dissemination techniques that capitalize on the statistical properties of the underlying social structures.

In summary, pervasive computing systems and technologies play an increasingly central role in the exploration of human dynamics and social networks, and vice-versa. Thus, understanding human behavior and developing data-driven techniques for conducting analyses can greatly facilitate their effective operation. To expand the impact of this work, the research community will have to investigate with priority the following problems:

- Address issues related to personal data use and privacy protection, first by revisiting the principles of underlying data ownership, and second, through the development of mechanisms that allow individuals to selectively reveal personal information in a controlled, verifiable and intuitive manner.
- Develop techniques for efficient analysis of very large scale data sets collected from real deployments of pervasive computing systems at scale, with large numbers of users, covering an extended geographic area and harvested through a variety of sensing modalities.
- Explore a wider variety of dynamics within realistic situations and patterns of use, progressing far beyond the scope of current experimentation!

3.2. Wearable computing⁵

Wearable Computing (WC) has an important role in observing/measuring human behaviors from the virtual world. The notion of WC emerged in the nineties⁶ when research groups around Sandy Pentland at MIT Media Lab, Dan Siewiorek at CMU, Steve Feiner at Columbia University and some others explored Mark Weiser's vision of "a disappearing computer,

⁵ By Gerhard Tröster.

⁶ The initiation of the *International Symposium on Wearable Computers ISWC* at MIT Boston in 1997 marks a milestone in the scientific perception of Wearable Computing; <http://www.iswc.net>.

(· · ·) woven into the fabric of our everyday lives” [38]. The mission of a wearable computer – originally formulated in 1997 – sketches a personal assistant, which is continuously available, unobtrusively integrated in our daily outfit, enabling extended perception, providing context-aware functionality and proactive support in information processing. Steve Mann, one of the early pioneers, stated the term of Humanistic Intelligence as “a new family of applications (· · ·), in which the body-worn apparatus augments and mediates the human senses” [39].

Looking back, WC has created or at least influenced several research fields, which have been established over the years, e.g. smart textiles, body sensor networks, low power and mobile computing, context recognition and context awareness, or human computer interaction.⁷

The maturity of WC can be demonstrated by several projects showing its main application domains. Supporting healthcare emerged as one of the first WC applications. Monitoring the wearer’s vital signs promises improved treatment and reduction of medical costs [40]. Many projects like the EU FP6 MyHeart⁸ are aiming at a preventive lifestyle and early diagnosis, by focusing on the integration of healthcare more seamlessly into everyday life.

The concept of a personal assistant goes with sport and leisure activities: the wearable system operates as a sport trainer who monitors the motions in combination with the performance and gives online feedback as investigated in the Swim Master [41].

As early as 1998, the relevance of WC in diagnosis and maintenance has been demonstrated: using a wearable system, connected to remote helpdesk and expertise center, the efficiency and accuracy of maintenance work could be improved [42].

Since in the early stage of WC, no computing devices were available suitable for mobile sensing and computing experiments, several groups designed and built their machines, such as VuMAN (CMU) [43], MIThril (MIT) [44] or the QBIC (ETH) [45]. That wearable landscape and therewith the focus of WC has changed with the emergence of the smartphone: this device accomplishes many of the visions in wearable computing; it is unobtrusive, always connected, and moreover, it offers the sensory platform necessary for the recognition of context awareness. For example, eight sensor modalities are implemented in the iPhone4: accelerometer, GPS, ambient light, microphones, proximity, cameras, compass and gyroscope. The connection to Internet services like local weather forecast and location expands the sensorial capabilities. But additional sensor modalities close to the human body are essential to adapt the smartphone to the individual needs of the users, such as monitoring the vital parameters using ECG electrodes.

Many future research activities in WC will focus on the exploitation of the smartphone-platform proliferation (from 240 million sold devices in 2010 to around 620 million in 2015). Three main research topics in WC are considered here, the applicability for daily use, the integration in the daily outfit and the exploration of new application fields.

Despite the improved handling of the mobile devices, the user interface claims attention and distracts the user, e.g. for scrolling the touch panel. Context aware user support combined with a more intuitive control reduces the cognitive load in using mobile phones. WC provides the sensors and tools to trace the individual context of the user. Communication through gestures, touch and body language are common in social interactions. Such communication interfaces would benefit users of mobile devices by providing them with a more intuitive and natural feeling of interaction.

Furthermore, the continuous adaptation to the pervasive computing environment without flooding the user with information requires a context driven selection of data [46]. Future wearable and pervasive systems cannot rely on statically deployed sensors and services, rather on dynamically varying sensor setups, close to the human body or in the environment around it. On the contrary, future mobile systems have to deal with the opportunistically discovered data sources characterized by varying and cooperative sensor arrangements. This opportunistic approach is going to be established as a broad research topic also in WC.⁹

Smart textile is a further aspect of WC. A garment with integrated conductive fibres and sensors has already been presented in 1997 [47]. But except for a few applications in sport and the military, smart textile is stuck at the prototype level mainly because the way in the established production chain has not been paved over the years. We expect that using commercial weaving machines [48] could offer the potential for adopting unobtrusively textile sensors and actuators into our daily outfit.

New application domains in ubiquitous computing rely on the support by WC systems. Sensing human behavior continuously allows the monitoring and exploring social group dynamics [49], establishing the emerging field of computational social science [50].

Looking into the future, Wearable Computing will remain indispensable to empower people in that ubiquitous computing world ensuring personalized mobile ecosystems. But several research challenges have to be met. The proliferation of smart textiles and clothes requires scaled-up manufacturing processes to exploit economy of scale effects. Technical textiles in cars will probably pioneer the large-scale production of smart fabrics also utilizable for smart clothes. Despite valuable achievements in on-body context recognition over the past years, approved applications are still limited to a few domains,

⁷ The influence can be measured by analyzing the publication record. Starting with a few publications in 1995 the scientific relevance is documented by more than 2800 publications using the term ‘wearable computing’ in the title or topic according to the ISI web of Knowledge (Oct 2010). Several conferences like the *ACM International Conference on Ubiquitous Computing (UbiComp)* and the *IEEE International Conference on Pervasive Computing and communications (PerCom)* have extended the momentum initiated by WC.

⁸ <http://www.hitech-projects.com/euprojects/myheart/>.

⁹ <http://opportunity-project.eu/>.

e.g. sports and wellness. To open up further domains more annotated data sets and models have to be developed to portray the various facets of life [51]. In addition to opportunistic reasoning approaches, web based data gathering methods are supportive in fusing heterogeneous sensor modalities and in the automatic annotation of data streams.

4. Cyber world paradigms enabled by physical world interactions: Opportunistic Networking and Computing¹⁰

Exploiting human relationships in the virtual world is adding new dimensions to pervasive computing and communications paradigms. Specifically, this gives rise to the concept of the Opportunistic Networking and Computing (ONC) paradigm.

Opportunistic networking can be viewed as a natural evolution of ad hoc networking and the original concept of disruption tolerant networking (DTN). While ad hoc networking flourished as a research area, its applications are limited to a few in the military. The DTN paradigm uses the ‘receive, store and forward’ mode for data transfers through gateways that enable communication across temporarily disconnected Internet clouds. Opportunistic networking, on the other hand, assumes a much more fluid networking environment, and exploits mainly opportunistic contacts between pairs of devices, distributed in space and time, even in the absence of Internet connectivity. Unlike the ad hoc networking paradigm, opportunistic networking uses mobility to enhance data distribution through an increased number of contacts.

Opportunistic networks represent the first attempt to close the gap between human and network behaviors, by taking a user-centric approach to networking and exploiting user (nodes) mobility as an opportunity (rather than a challenge) to improve data forwarding [6].

In opportunistic networking, users exploit each others’ resources in terms of mobility, bandwidth and storage, in order to support communication. Opportunistic computing brings this concept one step further by exploiting *appropriate, but opportunistically available* resources to support user application tasks. In the opportunistic computing vision, applications running on a given user’s device can exploit resources (bandwidth, computation, content, sensors, etc.) available either on other users’ devices or in general any computing/communication resources available in the environment.

Billions of smart and powerful cell phones, coupled with user mobility and human social nature provide the impetus to the ONC paradigm that utilizes state of the art devices (e.g., smartphones, sensory devices, onboard computers in vehicles) and communication channels (e.g., WiFi, Bluetooth, WiMax, or LTE). ONC performs typical distributed computing tasks, paving the way for a plethora of multitasking applications. It also supports *complex distributed applications*, possibly under *severely challenged conditions*, where *connectivity is a scarce resource*.

Such distributed tasks as content acquisition and dissemination, collaboration, parallel task execution, remote task execution, load balancing, access to shared resources and others, can be supported in purely opportunistic environments. Utilizing the basic tasks and available resources, new applications can be developed to enrich everyday life, monitor the environment, manage natural resources, provide entertainment, and so on. Indeed, the potential of ONC is mind boggling, as this is a pervasive, versatile, compelling, enormous opportunity that has not been tapped as yet. Furthermore, human mobility and social interactions can be utilized to facilitate efficient content distribution and distributed task execution.

Opportunistic sensing is another area of direct impact. User mobility and their social activities have the potential to facilitate opportunistic sensing of the physical environment with the help of devices equipped with sensors [52–54]. For example, users’ devices can be used to compute average temperature/humidity in public places to control heating/cooling devices; a jogger’s device can report an excessive pollen area to inform friends prone to allergies.

ONC is a notable example of the general concept of convergence between the physical and the cyber worlds. On the one hand, ONC exploits structures of social relationships present in the physical world to optimize the performance of service infrastructure in the cyber world. On the other hand, it organizes the resources available in the cyber world (which possibly abstract resources of the physical world), in order to provide enhanced functionality to the users in the physical world.

ONC enables interesting applications that have already started to emerge in the research landscape – pervasive healthcare, mobile social networking, participatory sensing, transportation, pervasive sustainability, entertainment, crisis management, the military and others. All these applications exploit the convergence and tight interactions between physical and cyber world elements.

As an example, the MetroSense project at Dartmouth College (<http://metrosense.cs.dartmouth.edu/>) investigates the concept of people-centric sensing [55]. They consider an environment similar to the one featuring opportunistic computing, and exploit mobile devices’ resources and sensory readings to infer user activities and provide improved mobile social networking services to facilitate social interactions among users.

For pervasive healthcare, a combination of participatory sensing, opportunistic contacts and services and social networking can be used to implement continuous monitoring and care. Intelligent transportation systems may also benefit a lot from the ONC concepts. They naturally fit vehicular scenarios, in which communication is necessarily sporadic and groups of vehicles join and split forming dynamic networks all the time, in addition to exploiting opportunistic connections to roadside units. ONC techniques can be used either for safety-oriented applications (e.g., as in the SAFESPOT EU Project, <http://www.safespot-eu.org/>) or for infotainment applications, as well as for supporting advanced multi-modal mobility services (such as opportunistic ride sharing services, for example).

¹⁰ By Mohan J. Kumar and Andrea Passarella.

ONC can also be used in crisis management scenarios, to support the quick establishment of networking and computing services in emergency scenarios when the primary infrastructures might not be available due to disruption or congestion. In this case, ONC can be used as a sort of “overlay” technology to opportunistically exploit any resource still available in the aftermath of accidents to provide, e.g., communication and information services to citizens and rescuers [56]. More details and examples related to these areas can be found in [9].

Despite the enormous potential, there are several challenges that need to be addressed to bring the ONC paradigm into everyday life. At the outset, we highlight five major “grand” challenges unique to ONC:

1. Understanding of the interplay between human social interactions and ONC solutions.
2. Designing scalable solutions to handle large number of devices and the huge amount of content shared and generated by users (and their devices).
3. Utilization of opportunistically encountered resources efficiently and securely.
4. Motivating users (or their mobile devices) to participate in packet forwarding and participatory sensing operations through new incentivizing mechanisms.
5. Efficient and effective management of information flow and service composition in the face of inevitable data/service duplications and resource limitations.

While a crisp separation is difficult, the first two challenges mainly touch upon networking aspects, while the last three also require considering distributed computing issues. For the sake of presentation, they are separately elaborated in Sections 4.1 and 4.2.

4.1. Networking challenges of ONC

In its initial phase, research in opportunistic networking was naturally focused on routing and mobility modelling [6], as the two key issues to be addressed first. These efforts have highlighted the importance of integrating information about the social behavior of users in the design of ONC solutions. It is now clear that “sociality” helps ONC solutions, as social information nicely complements unstable topological information to design ONC protocols. However, a *more profound understanding of the interplay between the user’s social behavior and performance of networking solutions exploiting social information* is a key challenge still to be addressed. Models describing the behavior of social-aware routing protocols [57,58] could shed light on which parameters of the user social behavior determine the performance of routing in opportunistic networks. Furthermore, while the capacity of opportunistic networks using naïve protocols [37] has been investigated, the extent of the impact of social behavior on capacity, and on the trade-off between capacity and resource consumption, needs further research. Last but not least, ONC technologies allow supporting new types of social networks – Ephemeral Social Networks, which are created dynamically by users sharing a common interest possibly for a limited amount of time in a well-defined physical space (e.g., tourists visiting a city). ONC technologies should natively exploit and support the existence of Ephemeral Social Networks. A key challenge from this standpoint is therefore how to recognize and adapt to dynamic social networks that may “appear” among groups of users, possibly even for limited amounts of time.

A further set of challenges for opportunistic networks research stems from the vision of the Future Internet as a *content-centric network* [59,60]. This perspective is particularly suitable for ONC scenarios, as receiving content of interest from possibly unknown peers is definitely a clear use case for ONC users. A content-centric view to ONC poses huge scalability challenges. The amount of content possibly generated by users can be huge, and can only be expected to grow exponentially. Managing such massive amounts of content in a network poses extremely exciting challenges: *how should content be indexed? how should it be replicated? how should it be advertised and finally retrieved? how can all of these functions be designed in ONC environments without saturating the network and device resources?* These questions call for decentralized content management algorithms that can assess the relevance of content for particular users, replicate content only where required, leverage on distributed storage capacities, etc. All of these questions are – as of now – still largely unexplored. Some of these issues are at the core of the EU RECOGNITION project’s activities, <http://www.recognition-project.eu/>.

Finally, two additional challenges are worth mentioning. Opportunistic networks rely on user cooperation in order to deploy networking services. This is actually a multi-faceted and pervasive challenge that is discussed more in detail in Section 4.2.

Furthermore, most of the research in ONC has considered networks completely isolated from the rest of the Internet. However, opportunistic networks should be seen as part of a pervasive Internet, comprising heterogeneous wired and mobile networks. It is still largely unexplored *how opportunistic networks can be integrated into and benefit from the existence of “better connected” parts of the global pervasive Internet.*

4.2. Distributed computing challenges of ONC

Traditional distributed computing relies on continuously connected and reachable computing resources [61]. A major challenge in opportunistic environments is to deal with the absence of continuous connectivity. As with networking, the key is to exploit user/device mobility and social networking to facilitate distributed computing operations. Service provisioning, a key problem in distributed computing, can be tackled by manifesting each resource as a basic *service component*. An application level service can be composed by stitching several basic services together, exploiting available

resources encountered opportunistically. As resource availability can be rather scarce in opportunistic networks, a first key challenge is *how to use them optimally without leading to saturation*. This is a well-known issue also from a networking standpoint, as aggressive packet forwarding has the potential to saturate the network with packet overload, thus leading to inefficient utilization of resources and multiple copies of information. Likewise, spawning of requests for service executions during every contact will lead to unnecessary usage of bandwidth, computation and battery energy. Novel mechanisms at the middleware level are required to ensure efficient flow of information in the network, to achieve optimal utilization of resources. Novel query processing mechanisms are needed to process queries issued during opportunistic contacts and generate query plans to in an opportunistically distributed environment characterized by short-lived content. It will be critical to understand patterns of mobility, social interactions, user interests and trust relationships, and then exploit such knowledge for better resource management and service provisioning. To accomplish the above, analytical modeling, large-scale simulation and extensive experiments are required. With respect to the latter, experiments have to be carried out involving a large number of participants (and their devices) from different walks of life over significant periods of time in different environments (e.g., students at a university over a semester, customers at a mall over a week, spectators during a game in a stadium).

Service provisioning in opportunistic environments is accomplished by information exchange between pairs of nodes upon encountering. Another challenge worth mentioning is the *optimal usage of the limited resources available during such contacts*. Buffer memory, residual battery energy and transmission bandwidth of intermediate nodes are vital to opportunistic networks, and must be managed efficiently at the device level. These resources at all intermediate nodes are utilized for successful data transfers and service executions upon contacts. An *opportunistic contact* is defined by the product of the contact time and the available bandwidth. Contact time must be utilized to perform basic handshake operations, exchange of trust/privacy information as well as the packets for forwarding.

Furthermore, in service execution and composition scenarios the contacting nodes must also exchange service states and parameters. New protocols are necessary to ensure swift assessment of the neighboring node's trust and potential. It will be necessary to develop new mechanisms that use context and social profile as prior information of the network and adapt network schemes, including routing, buffering, and device discovery.

Secure exchange of data during opportunistic computing is critical, as often the devices in contact may be unknown to each other. Issues related to *trust, security and cooperation* have to be addressed (see [63,62,64] for initial work in this direction). Also in this case, an interesting direction is to automatically translate trust relationships between users in their real social network in trust relationships between their devices [65]. Participation of intermediate nodes is the key to opportunistic networking. New mechanisms are necessary to encourage users to allow utilization of their device resources by extraneous services and applications. As existing mechanisms [57,6] are limited to packet forwarding, novel schemes for content/resource sharing and the other instances of opportunistic computing are necessary. There is an urgent need to carry out investigations and develop novel "business models" for opportunistic networks to motivate user involvement.

The discussion carried out so far allows us to break down the grand challenges identified in the beginning at a more granular level. We thus highlight some key research topics in the area of ONC that need to be urgently addressed:

- Analytical models of social aware networking protocols to better understand impact of social parameters on network protocols;
- design of networking protocols supporting dynamic, transient social networks, such as Ephemeral Social Networks;
- solutions for data dissemination, search, and indexing;
- exploitation of infrastructure entities to improve the performance of ONC services;
- models and distributed algorithms for resource management; and
- cooperation enforcement and trust assessment.

5. Information management in the cyber–physical world

The convergence between the cyber and physical world generates new challenges for handling the information flowing between the cyber and physical worlds. First, huge amount of information will be generated by several sources both in the physical (e.g., by humans or environmental phenomena) and in the virtual world (by sensors, RFIDs, etc.), which need to be stored for future processing. Furthermore, in order to use this huge amount of heterogeneous data coming from a variety of sources, the quality of such data needs to be evaluated and taken into consideration before using it.

In Section 5.1 we discuss the data storage challenges associated with the emerging cyber–physical systems, while in Section 5.2 we discuss the *Quality of Information (QoI)*.

5.1. Data storage¹¹

The arrival of the cyber–physical system era places new data storage and analysis requirements on pervasive and mobile computing. First, the volume of data keeps increasing rapidly, for the reason that a lot more sensors, such as RFID

¹¹ By Lionel M. Ni.

readers, mobile phones, and GPS are becoming available for continuous data collection. As a result, industrial and scientific users already or will soon work with datasets of peta-scale. Second, the structure of the data is becoming more complex. Unlike web data, which can be modeled as (key, value) pairs, data collected from sensors are multi-dimensional time series with spatial attributes. Therefore, traditional data models such as relational tables are no longer applicable. Third, the complexity of calculations is also increasing. For example, multi-dimensional aggregation and pattern discovery are becoming prominent, which may result in operations with poor performance.

The main challenges of data storage and analysis in pervasive and mobile computing are discussed in the following.

- *Scalability*

Computations on multi-petabytes of data inevitably pose significant challenges. Even a relatively simple operation becomes complex when dealing with multi-petabyte data. For example, a single sequential scan requires less than a second to go through a terabyte, while over a year to finish through a petabyte (at 10 MB/s). Even optimization techniques, such as indexes would be problematic when scaling to petascale; a single sequential scan at 10 MB/s through an index (which tends to be multi-terabytes) on a petabyte of data would take more than two days. The most effective technique to tackle this problem is parallel programming, among which the two most famous paradigms are parallel DBMS [66] and MapReduce [67]. For the sake of cost effectiveness, both paradigms adopt the shared-nothing clusters with low-end commodity machines [68]. However, neither of them is perfect. For MapReduce, it scores well in scalability when dealing with key-value pairs; however, a lot of work is yet to be done before applying it to multi-dimensional spatio-temporal streaming data. For parallel DBMS, due to the heavy semantics and poor fault tolerance, it is difficult to perform efficiently on clusters with only hundreds of nodes, not to mention the large data centers today that usually have tens of thousands of machines.

- *Complexity*

First, the structure of the data tends to be complex [69]. For example, in a typical digital city system, data may be time-series with spatial attributes (e.g., moving taxis), or spatial objects with extent (e.g., rivers, buildings, roads, etc.). Multi-point correlations in space and time are therefore required, engendering complex processing. Although a plethora of efforts have been made, there is no standard algebra defined on spatio-temporal data. Besides, raw data collected from sensors are likely to be incomplete and inaccurate. Such data adds uncertain ranges and calculation overheads to queries, further increasing the complexity. Second, calculations are getting more complex as well [70]. For example, one prominent task in both industrial and scientific analysis is to recognize patterns in (monitored) object behaviors. This task can be used for many business purposes, such as advertising and promotions, predicting churn, finding fraud, and analyzing social networks. To this end, various calculation techniques are introduced, ranging from coordination transformations on raw data to advanced machine learning algorithms on the derived data.

- *Flexibility*

Most query loads are highly unpredictable, with up to 90% of queries being new [71]. In existing systems, however, the data organization schemes are optimized based on built-in assumptions. Although modern techniques (e.g., materialized views), to some extent, can tune the system with query logs, they are not adaptable enough to the variation of query loads. Thus, the system will perform well only when these assumptions hold, and thus lack of flexibility. To enable high flexibility, a system should be vitalized enough to sense the variations of query loads, predict the potential future queries, and reorganize data correspondingly ahead of time. In [69], a new paradigm called Data Vitalization was proposed to meet these requirements. Though Data Vitalization is promising, refinement works are needed before implementing it in a real system.

- *Reproducibility*

As both the scale and complexity of datasets are increasing, it is critical to be able to reproduce a computational procedure and its result [71,72]. Usually the data pre-processing, such as data clean and data calibration, is fulfilled by external systems. Therefore, it is necessary to take into account the provenance information from both outside and inside the datasets. Besides, simply tracking what happened in history is far from enough. Tasks including maintaining lineage of data and the ability to use various versions of them should be involved as well. Obviously, great challenges need to be overcome to enable perfect provenance. First, it is pointed out that in a typical large-scale dataset, only 10% of the disk space is used for storing “hot” data, while the remaining 90% space is “wasted” for storing provenance information. Second, a computational procedure and its results are affected by numerous factors, such as the hardware conditions, the operation system, and the compiler library. It would be very expensive or even impossible to reproduce data under a given configuration.

5.2. Quality of the information¹²

With the proliferation of devices, connectivity alternatives, applications, and services, a key challenge becomes how to *orchestrate* them horizontally to allow seamless interoperation of all these for the benefit of end-users. Typically referred to as the *Internet of Things* (IoT), an abundance of addressable computing intelligence, with sensing and possibly actuating

¹² By Chatschik Bisdikian. His work on QoI has been performed through participation in the International Technology Alliance sponsored by the US Army Research Laboratory and the UK Ministry of Defence.

capabilities, will be embedded in physical (hence, cyber–physical) objects to interact with their physical surroundings, collecting information from them, disseminating to whoever needs them, and smartly (re)acting upon their desires. In the true spirit of the Internet, the binding and the flow of information between information sources (e.g., sensors), information processors (e.g., fusion elements), and information users (e.g., decision makers and actuators) will be dynamic, cross-domain, and, of course, open. This is in contrast to the typical deployment of, say, current sensor-enabled applications where sensing assets and applications are “monolithically” deployed and scoped to a particular purpose.

Consider the following (not so unrealistic) scenario: Due to sudden and severe wind outbreaks, falling branches from trees along a main one-way street in a residential area block the street and cut electricity wires interrupting the power supply in the neighborhood and a nearby busy retail area. Unrelated to this situation, an ambulance is speeding to pick-up a severely ill person living in a house a little bit further down the same road, on the same block where the branches fell. Sensors (owned by the power utility company) detect the power supply disruption and trigger battery-backed cameras (owned by the local police department) in the area to take visuals of the situation. Upon confirming and assessing the severity of the situation, a team of experts from the utility operator is summoned to deal with the power supply disruption. At the same time, the transportation department is notified that fallen branches have closed the road which in turn triggers a notification to the city parks department (which maintains the trees in the area) and sanitation department calling them to action in clearing the road from the debris. Furthermore, the transportation department issues a search for emergency response activity in the area that could be affected by the blocked road. The search points to our speeding ambulance owned by a private EMT (*Emergency Medical Transport*) company and tracked by an onboard GPS tracker rushing down the blocked one-way street. Based on the traffic conditions in the area, collected from sensors (owned by the city transportation department) embedded in the city roads, and in consultation with the EMT company which shared the ambulance’s destination, a new alternative path is plotted. While the EMT ambulance dispatcher instructs the ambulance driver to change course, updated driving directions are drawn on an electronic map on the ambulance’s dashboard. While this is happening, the transportation department tweaks the traffic signage and flow in the area to clear the path for the ambulance to reach its destination safely and quickly from the wrong way of the one-way street.

The above scenario provides an example of what horizontal, seamless interoperation of devices and services on IoT could enable. The scenario underscores a number of challenges such as scalability of information management procedures, security, dynamic service composition, open cross-domain operations, etc., whose magnitude and extension of the IoT will exacerbate to a degree well above and beyond the levels that can practically be attained by the current state-of-the-art technology.

In orchestrating rich, dynamic information flows and service compositions to support the creation and operation of smart pervasive solutions in ad-hoc created situations over highly diversified environments, the *quality of (sensor-derived) information* (QoI) will play a very significant role. Broadly speaking, QoI relates to the ability to judge whether available information is fit-for-use (for a particular purpose) [73]; it can form the basis for searching, selecting, and binding to the most pertinent information providers. QoI relates to the selection and dissemination of desired information and goes beyond traditional quality of service (QoS), which pertains only to transporting a bit, a byte, or a packet from point A to point(s) B within desired network parameters (bandwidth, jitter, loss probability). Also, QoI goes beyond the resource-constrained operation of sensor networks and also beyond the calculation of merely the accuracy of, say, a localization algorithm. QoI builds and extends upon all of these jointly with a purpose of balancing the design, deployment and operation strategies of information dissemination and processing/management networks against the information needs of the applications that these networks are dynamically being brought together to support [74,75].

The area of information quality has long been studied in the areas of enterprise data management, dating back to the ’90s and the work of Richard Wang and his team [76,77], and more recently within the context of Web searches [78]. For the former, the focus has been in the development of data acquisition, storage, and manipulation processes that are cognizant of potential quality degradations that these processes may entail, while, for the latter, it has been the assessment of the quality of information collected from the Web and, in general, document searches. In both cases, quality was considered along multiple dimensions (in one count there were well over 100) a small subset of which includes: *accuracy*, *timeliness*, *confidence* (including attributes such as *trust*), *completeness*, *relevancy* and *usability*. However, as the above scenario showcases, the (often resource-constrained) operational alternatives and the plurality of contexts that cyber–physical systems will be dealing with will be significantly higher than anything considered thus far in these domains. As a result, for cyber–physical systems, QoI needs to be a core part of their design, deployment, and operation, i.e., a key part of their architectural philosophy.

To this end, QoI as a dedicated middleware component for sensor and actuator networks was hinted in [79] and considered as a management component for MAC protocol operation in [80]. More recently, [81] introduced a concrete instance of a QoI middleware component in the form of a QoI-aware admission control module for sensing tasks built around a QoI satisfaction index for sensing tasks and a runtime learning of the relationship linking resource allocation and the satisfaction index. Quality aspects in middleware systems for sensor networks were also considered in [82] but the emphasis there was closer to networking adaptations for QoS satisfaction rather than for QoI directly. The survey in [83] on context modeling and reasoning techniques included studies capturing and describing context specifically under uncertainty conditions, which could potentially seed more focused research on QoI.

Outside a middleware framework, QoI has been considered (directly or indirectly) in various point cases. For example, [84] considered collaborating sensors in a sensor field for increasing the information utility in an object-tracking

scenario. Ref. [85] considered devising quality-driven sensor allocation schemes built around Bayesian networks that are used to describe the sensing tasks, actions, and reasoning that sensors are called to support. Ref. [86] considered explicitly the multi-dimensional aspects of quality and studies linkages between energy consumption in duty-cycled sensor networks and accuracy and latency in tracking moving objects. Trade-offs between energy-consumption and detection accuracy in networks of binary sensors was also considered in [87], which also was the inspiration behind [88], which presented an early study of the interplay between sensor operation, application characteristics and detection accuracy for detectors of transient phenomena. Ref. [89] studied the trade-off between the cost of caching data vs. staleness for a number of data lookup policies and for various data changing behaviors.

Aside from the computational aspects of QoI attributes, [90] considers a data model for capturing and propagating quality attributes in data streams such as those generated from monitoring machine operations. SensorML is a modeling and XML-based encoding scheme that provides a framework for describing the observational characteristics (including data quality) of sensor systems [91]. Pertinent to the latter are lightweight information descriptors for pervasive applications such as the xW descriptors $4W$ [92] and its derivative $5WH$ [93]. They both pertain to the *what*, *where*, and *when* of the information but serve somewhat different objectives, the former as a concise means for representing contextual information while the latter as a means to describe information needs for consumers and information producing capabilities producers of information, see also [73]. Regarding the *where*, [94] studies the selection of information providers that are relevant in a spatial sense, in that they produce information that has quality and spatial coverage attributes that are close, in some sense, to those that an information consumer desires.

More than any other of their characteristics, such as resource-constrained operation, sensor networks are recognized for their *edge-of-the-network* responsibilities. They are deployed to support the information needs of the end-user applications. As mentioned before, sometimes these applications are an integral part of the deployed sensor networks and sometimes they are not. These networks and applications could have been designed, built, deployed, administered, managed and operated by entirely different organizations. With the emergence of ad-hoc, late-binding, on-demand information exchanges involving (myriads of originally) unaffiliated smart objects in IoT settings, or personal devices on-the-move in participatory sensing settings, it becomes clear that pervasive systems, applications, and services need to be built around higher-value principles that go beyond node-centric and data-centric designs. Instead, information-centric principles will have to guide designs where the impact of a piece of information on the action that uses it ought to be assessed and accounted for. Just as QoS serves as the operational and management centrepiece of distributed systems today, QoI with its many dimensions can serve as the operational and management centrepiece of the systems of tomorrow. For an early example of linking QoS with aspects of QoI in data collection systems, see [95]. This would require a holistic approach to research on QoI and its implication to system operation and impact on end users.

Dealing with the collection, distribution, processing and storage of an abundance of information generated by a very large number of highly diversified data sources that are distributed over large geographical areas and owned, deployed and operated by different agencies will require a new breed of highly scalable algorithms to be devised, new system architectures to be created, and new system designs to be implemented. While the cited references provide good starting points, they represent examples of only narrowly scoped realizations of the required holistic QoI designs. To attain the broad benefits of end-user-centric designs, it will be required to add emphasis to QoI research focusing on areas such as:

- The study of methodologies (such as mathematical, data, and semantic models) that succinctly capture and communicate an application's information needs and desired quality and, likewise, the information-producing capabilities and achievable quality with respect to quality attributes such as accuracy, latency, security, privacy, context, provenance, etc.
- The study of relationships between QoI obtained from information producing systems (such as sensor networks) and the operation of these systems such as sensor misbehaviors, resource constrains, network performance, trust, etc.
- The study of metrics that capture the value that information of given quality levels brings to applications that may consume it.
- The study of design and deployment choices, performance trade-offs, and runtime management of information-producing and processing networks to support agile, QoI-aware, dynamic, pervasive IoT solutions that span across multi-administrative domains.

6. Cyber-world security¹³

The proliferation of small and increasingly inexpensive computing devices into all spheres of everyday life, prompts a number of challenges related to their security and privacy, e.g., [96]. In this section, we discuss two topics in secure pervasive computing that (though clearly not ignored thus far) are in need of further progress and attention from the research community.

¹³ By Gene Tsudik.

6.1. Attestation and code update for embedded devices and networks

Embedded computing devices have been steadily increasing in both number and variety. They include stand-alone devices – such as sensors, actuators, small personal (e.g., medical) appliances and RFID tags – as well as various computer peripherals, such as: keyboards, mice, cameras and thumb drives/dongles. At least some of them can be found in almost every household and office.

At the same time, great strides have been made by malware producers: state-of-the-art viruses and worms are now capable of infecting and penetrating not only their usual targets (laptops and desktops) but also cell-phones, PDAs and all kinds of embedded devices. In other words, software miniaturization triggers very real threats to devices previously thought to be out of bounds for malware. This, in turn, prompts the need for secure attestation and code update techniques.

The problem of secure code attestation can be summarized as follows: *how can an embedded device (prover) efficiently and securely prove that it runs the intended software and nothing else?* Secure code update is a related issue that, in addition, involves the installation of new software either (or both) before or after attestation. Attestation and code update have been investigated in the last few years and a number of techniques have been devised, varying widely in terms of efficiency, security and environmental assumptions. Solutions generally fall into several classes:

- **Secure Hardware:** techniques that require the presence of an on-board secure hardware component (e.g., a Trusted Platform Module (TPM) [97,98]) to assist in attestation. Since secure hardware generally guarantees the secrecy and non-malleability of code running within, it can be used as a “root of trust” and leveraged to bootstrap the secure attestation procedure. The main disadvantages of secure hardware are its added cost, power consumption and real estate needed to house it. However, as long as the assumption that the adversary cannot modify any device hardware holds, carefully designed secure hardware-based techniques can offer provably secure attestation. On the other hand, some devices are well suited for secure hardware-based methods as they routinely include TPMs, e.g., cell-phones and photo/video cameras.
- **Pure Software:** techniques that rely on clever usage of specialized attestation code (either present on the device or downloaded in real time) to verify absence of extraneous code or data, i.e., potential malware. Another common feature of software-based techniques is their reliance on precise timing of the attestation procedure. The reasoning is that any malware must occupy extra space on the prover device and thus must necessarily spend extra time in moving around (and/or compressing) legitimate code in order to successfully complete attestation. Although such techniques are appreciably cheaper (as they require no additional components), they generally offer uncertain, or at least not provable, security. To wit, certain recent results have cast some doubt on the security of pure software-based attestation [99].
- **Hybrids:** techniques that assume something between secure hardware and pure software. This includes techniques based on the availability of read-only memory (ROM) housing small and immutable attestation code [100]. Unfortunately, current ROM-based approaches involve high communication costs and assume adversarial silence during attestation; hence they are limited to settings where the prover and the verifier are one hop apart.

The current state of affairs in embedded device attestation incentivizes further research. Many types of small and cheap devices cannot be outfitted with secure hardware. They also might be incapable of engaging in time- and bandwidth-consuming protocols needed by current hybrid methods, while, as mentioned above, software-based methods require precise timing and offer uncertain security. Furthermore, even in the domain of hardware-based techniques, there remains an open question: *what is the absolute minimum set of features needed from the TPM to support secure attestation?* Answering this question can help determine the lower bound on the cost of hardware-based techniques. An interesting treatise on the challenges posed by attestation can be found in [101]. Even if (or when) efficient, secure and low-cost attestation techniques materialize, the next research challenge is how to utilize them to perform attestation of many devices, e.g., an entire network of sensors or actuators.

6.2. Usable and secure configuration of pervasive devices

A major challenge in modern pervasive and ubiquitous computing stems from the usability of secure configuration and association of various devices. There are three main reasons for this challenge:

- First, the number of wireless devices is increasing rapidly and penetrating all layers of society. For example, while in the past they were owned by privileged few, cell-phones and their accessories (e.g., Bluetooth headsets) are common today throughout the entire world, permeating throughout all socio-economic strata. Consequently, their secure configuration is a task being faced by average non-specialist users.
- Second, the range of consumer-oriented wireless devices is broadening. Besides the usual laptops, PDAs and smart-phones, more and more everyday office and household items are acquiring a “wireless personality”. They include car keys, remote controls (as well as appliances they control), digital music players, cameras, alarms, stoves, refrigerators, medical aids (e.g., pacemakers, insulin monitors), utility meters, and a slew of other devices.
- Third, the actual processes of initial configuration of, and association among, wireless devices must be secure. It is not enough for them to result in secure access to, or channels among, devices. If these initial tasks themselves are not performed securely, all subsequent security is jeopardized or uncertain at best. Specifically, because devices

communicate wirelessly, there is no way for the human user to ascertain communication, i.e., to determine which devices are communicating. This triggers very realistic threats of so-called “Man-in-the-Middle” (MiTM) and “Evil Twin” (ET) attacks.¹⁴

The problem has been explored and more-or-less addressed in the context of interface-rich devices such as PDAs and cell-phones [102,103]. These tend to have multiple human-imperceptible wireless channels (e.g., WiFi and Bluetooth) as well as human-perceptible input (e.g., keypad/keyboard, microphone) and output (e.g., screen, speaker) interfaces. However, even with such relatively sophisticated devices, the problem of secure initial association is not as trivial as it might seem, as illustrated by numerous examples [104–106].

However, the problem gets gradually more difficult as devices become weaker, simpler and more specialized, while their user interfaces become more rudimentary. For example, wireless access points, sensors, wireless headsets and miniature music players present some obstacles for secure association and configuration since their user interfaces are typically limited to rudimentary audio (in the worst case, only output via beeps), simple visual (e.g., a single LED) or basic tactile input (e.g., one or two buttons). While techniques have been developed even for such very basic interfaces, they are awkward and not particularly loved by users [107–110].

There are three areas where further work is clearly needed. Consider a group of users, each with his/her personal wireless device (e.g., laptop, PDA or smart-phone), needing to set up a common secure channel for the purposes of, say, an impromptu conference. Lacking any prior secure state, these devices are unfamiliar to each other. Setting up a secure channel requires the participation of every device and every user. Thus, scalability becomes an issue, on top of all other challenges encountered in a simpler two-device setting. It turns out that only a few techniques workable in a two-device case are applicable to the group setting. (This is despite the fact that, in this context, devices tend to be more powerful and interface-rich.) While some preliminary results appear promising [111–113], much remains to be done in terms of both developing new techniques and conducting comparative usability studies.

Another related scenario is where one user needs to securely set up (configure and/or associate) a multitude of homogeneous wireless-capable devices, such as sensors. (These sensors might be: alarm system components, garden sprinklers, smoke/fire detectors, light switches, or electrical outlets.¹⁵) The difficulty of working with simple specialized devices is exacerbated by their number, i.e., scale and the fact that only one user is involved. Although mass secure initialization of sensors has been investigated [114,115], it involves either (or both) undue user burden or specialized equipment. It remains unclear whether better usability is attainable.

The final setting involves one user who has a home or office network composed of heterogeneous wireless devices. The user in question needs to introduce a new device into the existing secure wireless environment. Once again, this “enrolment” process must not be burdensome and, at the same time, be secure and amenable to a wide variety of devices. The problem is complicated by the fact that not all current devices might be “awake” or physically present during the new device’s enrolment. Also, because one user manages the entire network, it is reasonable to assume the presence of some sort of a “master device” owned and operated by that user. This prompts the usual concerns of having a single Trusted Third Party, including the consequences of its potential failure, loss, compromise or non-availability.

7. Discussions and conclusions

In this paper, we have explored new research challenges emerging from the convergence of cyber-and physical worlds. This convergence will be characterized by huge amount of information flowing from the physical world to the cyber world (and vice-versa), thus opening up great research opportunities for pervasive and mobile computing researchers. Information collected by sensors spread in the physical world is a fundamental tool for understanding the human behavior and socio/physical context in which humans operate. This information can be used for adapting the cyber world to human needs and possibly also affecting the physical world through actuators. The understanding of human behavior and social links constitutes the basis for the development of innovative computing and communication paradigms (e.g., opportunistic computing and networking) based on (mobile) social networks. As the information will be a key element in the CPW convergence, appropriate techniques are needed to evaluate the quality of the information, select the appropriate information sources, analyze them and store the information for easy access and retrieval.

The complexity of the cyber world, and the need to adapt its behavior to the human/social context, requires the developments of self organizing strategies for adapting the huge number of devices in the cyber world to the rapidly changing physical/social world. In addition, the proliferation of small and increasingly inexpensive computing devices into the physical world prompts a number of challenges related to their security and privacy.

The research challenges discussed in this manuscript represent a notable, but not exhaustive, list of research opportunities generated by the CPW convergence. Many other research challenges can be identified. For example, urban/participatory sensing techniques have a major role in understanding both the physical (e.g., human behavior, the

¹⁴ An MiTM attack occurs whenever an adversary electronically interposes itself between two legitimate entities and masquerades as one of them to the other (and vice versa). Whereas, an ET attack involves an adversary that simply pretends to be one of the devices.

¹⁵ Some of these examples correspond to items already on the market, while others are quite plausible and might soon become commonplace.

context in which humans operate, etc.) and the virtual world (e.g., virtual social structures, Internet of things, etc.) [116, 52]. Data sensing techniques produce very precious information (to understand the human and social context) that have confidentiality requirements and hence need to be carefully handled to preserve data privacy [53,11]. Advanced data mining techniques need to be developed/applied to extract meaningful information from sensed data, and complex-system techniques are required for modeling the human and social behaviors emerging from observed data. Device mobility (by expanding the Internet edges) is providing an additional spatial dimension to data acquisition and dissemination of content [7,9]. In addition, the socio-technical nature of the CPW convergence calls for novel and interdisciplinary research approaches mixing ICT (information and communication technologies) expertise with lessons learned from applied psychology, sociology, social anthropology, complexity science, etc.

Given the myriads of key research challenges arising out of the CPW convergence, let us conclude this article with three grand challenges that have been in-depth discussed in a workshop jointly sponsored by the US National Science Foundation and the European Commission, and collocated with IEEE PerCom 2010 [117].

1. *Understand and characterize the inter-relation between real-world social structures and online social networks for data dissemination in the cyber-physical world.*

The convergent cyber-physical world will be *information-* (or, *content-*) *centric*, and the users themselves will increasingly generate the information in a participatory fashion. In this scenario, where information is increasingly generated by the users and/or available on the users' devices, social networks will play a very important role in distributing content in the network on a massive scale by establishing overlays that link together the physical world (the users) and/or the cyber world (users' devices). Therefore, it is fundamental to acquire a better understanding of real world and online communities and their relationship. Of special interest are the opportunities to provide security, trust and privacy by exploiting the properties of human structures. Establishing trust and security for an interaction between *a priori* unknown peers is a challenging issue. However, social network structures offer a basis to enhance trust and security provision by capitalizing on existing social links.

2. *New paradigms for context determination in the cyber-physical world.*

Context modeling and characterization has been extensively investigated in the pervasive computing literature [118,119]. However, new methods must be investigated for gathering context data in social networks and more generally in the CPW [120]. Dealing with context uncertainty and predicting complex behavior in social networking systems is hard, due to increased ambiguities associated with decision making in scenarios or applications involving a group rather than individuals. New policies and guidelines are needed for gathering context data in social networks. Methodologies for opportunistic collection of context data are needed to utilize opportunistic sensing through users' devices. Furthermore, context determination in the presence of a large number of individuals is a new dimension to research in context-aware computing. Also, environmental or cyber context (i.e., what computing and communication resources are available in the neighborhood), and their relationship with social context is important.

3. *Privacy issues in participatory sensing of the cyber-physical world*

Participatory sensing has a major role in the CPW convergence and this creates major challenges for data privacy [53]. Combining sensed data from many mobile users can be a very powerful tool, but also risky since one can infer from sensed data not only location, but also other context such as how long one stays at one place, whom a person met, where one might be going, and other contextual and personal information. Therefore, we need privacy procedures, such as algorithms and protocols, to protect the sensed data against unwanted collection and distribution of personal information. However, it is worth noting that privacy is not a black or white proposition, and, like many other things, context influences its acceptable level. For example, a person may not wish to have all his/her steps sensed and reported to a certain organization without his control but might want to give up some privacy to a certain organization for receiving desired services in response. Furthermore we typically accept a lower level of privacy inside a social group.

In practice, privacy and performance have conflicting requirements and goals. It is critical to determine the balance to satisfy both points of view. Often, cultural backgrounds, psychology, geographic locations/barriers and regulatory policies play critical roles.

These three challenges are coherent with those identified in this paper, where we have in addition outlined the need for tackling issues related to autonomic behavior, opportunistic networking and computing, and quality of information (contextual or otherwise). Additional challenges related to pervasive computing at scale were discussed at the National Science Foundation's PeCS workshop [121].

References

- [1] D.J. Cook, S.K. Das, How smart are our environments? An updated look at the state of the art, *Pervasive and Mobile Computing* (Special Issue on Smart Environments) 3 (2) (2007) 53–73.
- [2] S.K. Das, N. Roy, A. Roy, Context-aware resource management in multi-inhabitant smart homes: a framework based on Nash H -learning, *Pervasive and Mobile Computing* (Special Issue on IEEE PerCom 2006 Selected Papers) 2 (4) (2006) 372–404.
- [3] D.J. Cook, S.K. Das, *Smart Environments: Technology, Protocols and Applications*, John Wiley, 2005.
- [4] F.-J. Wu, Y.-F. Kao, Y.-C. Tseng, From wireless sensor networks towards cyber-physical systems, *Pervasive and Mobile Computing* 7 (4) (2011) 397–413.
- [5] D.J. Cook, J.C. Augusto, V.R. Jakkula, Ambient intelligence: technologies, applications, and opportunities, *Pervasive and Mobile Computing* 5 (4) (2009) 277–298.

- [6] L. Pelusi, A. Passarella, M. Conti, Opportunistic networking: capitalize or not? Consistent data forwarding in disconnected mobile ad hoc networks, *IEEE Communications Magazine* (2006) 134–141.
- [7] M. Conti, S. Giordano, Martin May, A. Passarella, From opportunistic networks to opportunistic computing, *IEEE Communications Magazine* 48 (9) (2010) 126–139.
- [8] M. Conti, Special section on mobile opportunistic networking, *Pervasive and Mobile Computing* 7 (2) (2011) 159–222.
- [9] M. Conti, M. Kumar, Opportunities in opportunistic computing, *IEEE Computer* 43 (1) (2010) 42–50.
- [10] S.K. Das, K. Kant, N. Zhang, Handbook on Securing Cyber-Physical Critical Infrastructure, Morgan Kaufman, 2012.
- [11] N. Li, N. Zhang, S.K. Das, Preserving relation privacy in online social network data, *IEEE Internet Computing (Special Issue on Security and Privacy in Social Networks)* 15 (3) (2011).
- [12] S. Dobson, S. Denazis, A. Fernández, D. Gaiti, E. Gelenbe, F. Massacci, P. Nixon, F. Saffre, N. Schmidt, F. Zambonelli, A survey of autonomic communications, *ACM Transactions on Autonomous and Adaptive Systems* 1 (2) (2006) 223–259.
- [13] J. Kephart, D.M. Chess, The vision of autonomic computing, *IEEE Computer* 36 (1) (2003) 41–50.
- [14] O. Baboglu, et al., Design patterns from biology to distributed computing, *ACM Transaction on Autonomous and Adaptive Systems* 1 (1) (2006) 26–66.
- [15] M.C. Yuen, L.J. Chen, I. King, A survey of human computation systems, in: *Proceedings of the International Conference on Computational Science and Engineering*, Vancouver, CA, 2009, IEEE CS Press.
- [16] C. Bettini, et al., A survey of context modeling and reasoning techniques, *Pervasive and Mobile Computing* 6 (2) (2010) 161–180.
- [17] N. Biccocchi, et al., Self-organized data ecologies for situation-aware pervasive services: the knowledge networks approach, *IEEE Transactions on Systems, Man, and Cybernetics—Part A* 40 (4) (2010) 789–802.
- [18] B.H.C. Cheng, et al., Software engineering for self-adaptive systems: a research roadmap, in: *Self-Adaptive Software*, in: LNCS, vol. 5525, Springer-Verlag, 2009, pp. 1–16.
- [19] A. Halevy, P. Norvig, F. Pereira, The unreasonable effectiveness of data, *IEEE Intelligent Systems* 24 (2) (2009) 8–12.
- [20] F. Zambonelli, Self-management and the many facets of nonself, *IEEE Intelligent Systems* 21 (2) (2006) 50–58.
- [21] A. Di Ferdinando, et al., MyAds, a system for adaptive pervasive advertisement, *Pervasive and Mobile Computing* 5 (5) (2009) 385–401.
- [22] J.-S. Lee, B. Hoh, Dynamic pricing incentive for participatory sensing, *Pervasive and Mobile Computing* 6 (6) (2010) 693–708.
- [23] J. Hendler, N. Shadbolt, W. Hall, T. Berners-Lee, D. Weitzner, Web science: an interdisciplinary approach to understanding the web, *Communication of the ACM* 51 (7) (2008) 60–69.
- [24] L. Song, U. Deshpande, U.C. Kozat, D. Kotz, R. Jain, Predictability of WLAN mobility and its effects on bandwidth provisioning, *INFOCOM* (2006).
- [25] V. Kostakos, E. O’Neill, A. Penn, G. Roussos, D. Papadogkonas, Brief encounters: sensing, modeling and visualizing urban mobility and copresence networks, *ACM Transactions Computer and Human Interaction* 17 (1) (2010).
- [26] A. Angus, D. Papadogkonas, G. Papamarkos, G. Roussos, G. Lane, K. Martin, N. West, S. Thelwall, Z. Sujon, R. Silverstone, Urban social tapestries, *IEEE Pervasive Computing (PERVASIVE)* 7 (4) (2008) 44–51.
- [27] F. Calabrese, J. Reades, C. Ratti, Eigenplaces: segmenting space through digital signatures, *IEEE Pervasive Computing (PERVASIVE)* 9 (2010) 78–84.
- [28] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, J. Scott, Impact of human mobility on opportunistic forwarding algorithms, *IEEE Transactions of Mobile Computing* 6 (6) (2007) 606–620.
- [29] C. Boldrini, A. Passarella, HCMM: modelling spatial and temporal properties of human mobility driven by users’ social relationships, *Computer Communications (COMCOM)* 33 (9) (2010) 1056–1074.
- [30] M.C. Gonzalez, C.A. Hidalgo R. A.-L. Barabási, Understanding individual human mobility patterns, *Nature* 453 (2008) 779–782.
- [31] G. Roussos, M. Musolesi, G.D. Magoulas, Human behavior in ubiquitous environments, *Modeling of Human Mobility Patterns* 6 (4) (2010) 399–496.
- [32] D. Lazer, A. Pentland, L. Adamic, S. Aral, A.-L. Barabási, D. Brewer, N. Christakis, et al., Computational social science, *Science* 323 (5915) (2009) 721–723.
- [33] P. Norvig, Practice makes perfect: how billions of examples lead to better models, *ETech Conference*, March 3–9, San Diego, CA, USA 2008.
- [34] J. Linn, C. Dyer, Data-intensive text processing with mapreduce, in: *Synthesis Lectures on Human Language Technologies*, Morgan and Claypool Publishers, 2010.
- [35] A. Narayanan, V. Shmatikov, Robust de-anonymization of large sparse datasets, in: *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, SP’08, IEEE Computer Society, Washington, DC, USA, 111–125 2008.
- [36] D. Quercia, N. Lathia, F. Calabrese, G. Di Lorenzo, J. Crowcroft, Recommending social events from mobile phone location data, 10th IEEE International Conference on Data Mining, ICDM 10, Sydney, Australia 2010.
- [37] M. Grossglauser, D. Tse, Mobility increases the capacity of ad-hoc wireless networks, *IEEE/ACM Transactions on Networking* 10 (4) (2002).
- [38] M. Weiser, The computer for the 21st century, *Scientific American* 265 (3) (1991) 66–75.
- [39] S. Mann, Wearable computing: toward humanistic intelligence, *IEEE Intelligent Systems (May/June)* (2001) 10–15.
- [40] A. Pentland, Healthwear: medical technology becomes wearable, *IEEE Computer (May)* (2004) 42–49.
- [41] M. Bächlin, K. Förster, G. Tröster, A wearable assistant for swimmer, *Proceedings of Ubicomp* (2009) 215–219.
- [42] D. Siewiorek, A. Smailagic, L. Bass, J. Siegel, R. Martin, B. Bennington, Adtranz, A mobile computing system for maintenance and collaboration, *Proceedings of ISWC* (1998) 25–32.
- [43] A. Smailagic, D. Siewiorek, System level design as applied to CMU wearable computers, *Journal of VLSI Signal Processing* 21 (3) (1999) 251–263.
- [44] R. DeVaul, M. Sung, J. Gips, A. Pentland, MITHril 2003: applications and architecture, *Seventh IEEE International Symposium on Wearable Computers*, ISWC’03, 2003, p.4.
- [45] O. Amft, M. Lauffer, S. Ossevoort, F. Macaluso, P. Lukowicz, G. Tröster, Design of the QBIC wearable computing platform, in: *Proceedings of the 15th IEEE International Conference on Application specific Systems, Architectures and Processors*, ASAP, Sept 2004.
- [46] P. Lukowicz, O. Amft, D. Roggen, J. Cheng, On-body sensing: from gesture-based input to activity driven interaction, *IEEE Computer (Oct.)* (2010) 92–96.
- [47] E.J. Ling, S. Jaharaman, S. Park, R. Rajamanickam, R. Eisler, G. Burghart, T. McKee, A Sensate liner for personal monitoring applications, *Proceedings of ISWC* (1997) 98–107.
- [48] K. Cherenack, C. Zysset, T. Kinkeldei, N. Münzenrieder, G. Tröster, Woven electronic fibers with sensing and display functions for smart textiles, *Advanced Materials* 22 (2010) 5178–5182.
- [49] Wen Dong, Alex Pentland, Quantifying group problem solving with stochastic analysis, in: *Proc. ACM ICMI-MLMI 2010*, Beijing, China, November 8–10, 2010.
- [50] D. Lazer, et al., Computational social science, *Science* 323 (2009) 721–723.
- [51] D. Roggen, A. Calatroni, M. Rossi, T. Holleczeck, K. Förster, P. Lukowicz, D. Bannach, G. Pirkel, A. Ferscha, J. Doppler, C. Holzmann, M. Kurz, G. Holl, R. Chavarriaga, M. Creatura, J. del, R. Millan, Collecting complex activity data sets in highly rich networked sensor environments, in: *7th International Conference on Networked Sensing Systems*, IEEE Press, 2010, pp. 233–240.
- [52] H. Lu, N.D. Lane, S.B. Eisenman, A.T. Campbell, Bubble-sensing: binding sensing tasks to the physical world, *Pervasive and Mobile Computing* 6 (1) (2010) 58–71.
- [53] M. Shin, C. Cornelius, D. Peebles, A. Kapadia, D. Kotz, N. Triandopoulos, AnonySense: a system for anonymous opportunistic sensing, *Pervasive and Mobile Computing* 7 (1) (2011) 16–30.
- [54] J.-S. Lee, B. Hoh, Dynamic pricing incentive for participatory sensing, *Pervasive and Mobile Computing* 6 (6) (2010) 693–708.
- [55] A. Campbell, S.B. Eisenman, N.D. Lane, E. Miluzzo, R.A. Peterson, H. Lu, X. Zheng, M. Musolesi, K. Fodor, G.-S. Ahn, The rise of people-centric sensing, *IEEE Internet Computing* 12 (4) (2008) 12–21.
- [56] R. Bruno, M. Conti, A. Passarella, Opportunistic networking overlays for ICT services in crisis, in: *Proceedings of the 5th International Conference on Information Systems for Crisis Response and Management*, ISCRAM 2008, Washington, DC, USA, May 4–7, 2008.

- [57] P. Hui, J. Crowcroft, E. Yoneki, BUBBLE Rap: social based forwarding in delay tolerant networks, in: 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), Hong Kong, May, 2008.
- [58] E.M. Daly, M. Haahr, Social network analysis for information flow in disconnected delay-tolerant manets, *IEEE Transactions in Mobile Computing* (May) (2009).
- [59] T. Koponen, M. Chawla, B. Chun, A. Ermolinskiy, K.H. Kim, S. Shenker, I. Stoica, A data-oriented (and beyond) network architecture, *SIGCOMM Computational Communication Review* 37 (4) (2007) 181–192.
- [60] V. Jacobson, D.K. Smetters, J.D. Thornton, M.F. Plass, N.H. Briggs, R.L. Braynard, Networking named content, in: Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies Rome, Italy, December 01–04, 2009, CoNEXT '09, 2009.
- [61] G. Coulouris, J. Dollimore, T. Kindberg, *Distributed Systems: Concepts and Design*, Addison Wesley, ISBN: 0201-619-180, June 2005.
- [62] N. Li, S.K. Das, A trust-based framework for data forwarding in opportunistic networks, *Ad Hoc Networks*, in: M.C. Vuran, W. Heinzelman, et al. (Guest Eds.), Special Issue on Wireless Communication in Challenged Environments 2011 (doi: 10.1016/j.adhoc.2011.01.018).
- [63] A. Shikfa, M. Önen, R. Molva, Privacy and confidentiality in context-based and epidemic forwarding, in: *Computer Communications*, Elsevier, April 2010.
- [64] N. Li, S.K. Das, RADON: reputation-assisted data forwarding in opportunistic networks, in: Proceedings of 2nd ACM International Workshop on Mobile Opportunistic Networking (MobiOpp), Pisa, Italy, Feb 2010.
- [65] S.M. Allen, G. Colombo, R.M. Whitaker, Cooperation through self similar social networks, *ACM Transactions of Autonomous Adaption of System* 5 (1) (2010) 1–29.
- [66] D. DeWitt, J. Gray, Parallel database systems: the future of high performance database systems, *Communications of the ACM* 35 (6) (1992) 85–98.
- [67] J. Dean, S. Ghemawat, MapReduce: simplified data processing on large clusters, in: OSDI'04: Proceedings of the 6th Conference on Symposium on Operating Systems Design & Implementation. San Francisco, CA: USENIX Association, 2004, pp. 1–13.
- [68] L. Barroso, J. Dean, U. Holzle, Web search for a planet: the Google cluster architecture, *IEEE Micro* 23 (2) (2003) 22–28.
- [69] Z. Xiong, W. Luo, L. Chen, L.M. Ni, Data vitalization: a new paradigm for large-scale dataset analysis, in: The IEEE Sixteenth International Conference on Parallel and Distributed Systems, 2010.
- [70] R.H. Güting, H. Ralf, M. Schneider, Moving Objects Databases, Morgan Kaufmann, 2005.
- [71] B. Jacek, K.-T. Lim, Report from the 2nd workshop on extremely large databases, *Data Science Journal* 7 (November) (2008) 196–208.
- [72] Y.L. Simmhan, B. Plale, D. Gannon, A survey of data provenance in e-science, *ACM SIGMOD Record* 34 (3) (2005) 31–36.
- [73] C. Bisdikian, L.M. Kaplan, M.B. Srivastava, D.J. Thornley, D. Verma, R.I. Young, Building principles for a quality of information specification for sensor information, in: 12th International Conference on Information Fusion, FUSION 2009, Seattle, WA, USA, July 6–9 2009.
- [74] N. Roy, A. Misra, C. Julien, S.K. Das, J. Biswas, An energy efficient quality adaptive multi-modal sensor framework for context recognition, in: Proceedings of the 11th IEEE International Conference on Pervasive Computing and Communications (PerCom), Seattle, WA, pp. 63–75, Mar 2011. (Best Paper candidate).
- [75] N. Roy, S.K. Das, C. Julien, Resource-optimized quality-assured ambiguous context mediation framework in pervasive environments, *IEEE Transactions on Mobile Computing*, 2011, <http://doi.ieeecomputersociety.org/10.1109/TMC.2011.20>.
- [76] R.Y. Wang, D.M. Strong, Beyond accuracy: what data quality means to data consumers, *Journal of Management Information Systems* (1996).
- [77] Y. Wand, R.Y. Wang, Anchoring data quality dimensions in ontological foundations, *Communications of ACM* (Nov.) (1996).
- [78] S.-A. Knight, J. Burn, Developing a framework for assessing information quality on the world wide web, in: Information Science & Information Technology Education Joint Conference (InSITE), Flagstaff, AZ, USA, June 16–19, 2005.
- [79] J.W. Branch, J.S. Davis, D.M. Sow, C. Bisdikian, Sentire: a framework for building middleware for sensor and actuator networks, in: 1st Int'l Workshop on Sensor Networks and Systems for Pervasive Computing, PerSeNS 2005, part of IEEE PerCom'05, Kauai Island, HI, USA, Mar. 8–12, 2005.
- [80] I. Hwang, Q. Han, Archana Misra, MASTAQ: a middleware architecture for sensor applications with statistical quality constraints, in: 1st Int'l Workshop on Sensor Networks and Systems for Pervasive Computing, PerSeNS 2005, part of IEEE PerCom'05, Kauai Island, HI, USA, Mar. 8–12, 2005.
- [81] C.H. Liu, C. Bisdikian, J.W. Branch, K.K. Leung, QoI-aware wireless sensor network management for dynamic multi-task operations, in: IEEE SECON, Boston, MA, USA, June 21–25, 2010.
- [82] H.S. Carvalho, W.B. Heinzelman, A.L. Murphy, C.J.N. Coelho, A general data fusion architecture, in: Proc. 6th ISFS Int'l Conf. Of Information Fusion, FUSION 2003, Cairns, Queensland, Australia, July 8–11, 2003.
- [83] C. Bettini, O. Brdiczka, K. Henricksen, J. Indulski, D. Nicklas, A. Ranganathan, D. Riboni, A survey of context modelling and reasoning techniques, *Pervasive and Mobile Computing* (2010).
- [84] F. Zhao, J. Shin, J. Reich, Information-driven dynamic sensor collaboration for tracking applications, *IEEE Signal Processing Magazine* (2002).
- [85] A. Tolstikov, C.-K. Tham, W. Xiao, J. Biswas, Information quality mapping in resource-constrained multi-modal data fusion system over wireless sensor network with losses, in: 6th Int'l Conf. on Information, Communications & Signal Processing, ICICS 2007, Singapore, Dec. 10–13, 2007.
- [86] S. Zahedi, M.B. Srivastava, C. Bisdikian, L.M. Kaplan, Quality tradeoffs in object tracking with duty-cycled sensor networks, in: 31st IEEE Real-Time Systems Symposium, RTSS 2010, San Diego, CA, USA, Nov. 30-Dec. 3, 2010.
- [87] L. Yu, L. Yuan, G. Qu, A. Ephremides, Energy-driven detection scheme with guaranteed accuracy, in: Information Processing in Sensor Networks, IPSN 2006, Nashville, TN, USA, April 19–21, 2006.
- [88] C. Bisdikian, On sensor sampling and quality of information: a starting point, in: 3rd Int'l Workshop on Sensor Networks and Systems for Pervasive Computing, PerSeNS 2007, White Plains, NY, USA, March 19–23, 2007.
- [89] D.J. Yates, E.M. Nahum, J.F. Kurose, P. Shenoy, Data quality and query cost in pervasive sensing systems, *Pervasive and Mobile Computing* (2008).
- [90] A. Klein, W. Lehner, Representing data quality in sensor data streaming environments, *ACM Journal of Data and Information Quality* 1 (2) (2009).
- [91] Open Geospatial Consortium, Sensor Model Language, <http://www.opengeospatial.org/standards/sensorml>.
- [92] G. Castelli, A. Rosi, M. Mamei, F. Zambonelli, A simple model and infrastructure for context-aware browsing of the world, in: 5th IEEE Int'l Conf. on Pervasive Computing and Communications, PerCom 2007, White Plain, New York, USA, March 19–23, 2007.
- [93] C. Bisdikian, J. Branch, K.K. Leung, R.I. Young, A letter soup for the quality of information in sensor networks, in: 1st Information Quality and Quality of Service Workshop, IQ2S'09, part of IEEE PerCom'09, Galveston, TX, USA, March 9–13, 2009.
- [94] G. Tychogiorgos, C. Bisdikian, Selecting relevant sensor providers for meeting 'your' quality information needs, in: 12th Int'l Conf. on Mobile Data Management (MDM 2011), Luleå, Sweden, June 6–9, 2011.
- [95] Q. Han, D. Hakkarinen, P. Boonma, J. Suzuki, Quality-aware sensor data collection, *International Journal of Sensor Networks* 7 (3) (2010).
- [96] W. Zhang, Y. Liu, S.K. Das, P. De, Secure data aggregation in wireless sensor networks: a watermark based authentication supportive approach, *Pervasive and Mobile Computing* 4 (5) (2008) 658–680.
- [97] Trusted Computing Group, Specifications, www.trustedcomputinggroup.org.
- [98] R. Sailer, X. Zhang, T. Jaeger, L. van Doorn, Design and implementation of a TCG-based integrity measurement architecture, *USENIX Security Symposium*, 2004.
- [99] C. Castelluccia, A. Francillon, D. Perito, C. Soriente, On the difficulty of software-based attestation of embedded devices, *ACM Conference on Computer and Communications Security, CCS'09*.
- [100] D. Perito, G. Tsudik, Secure code update for embedded devices via proofs of secure erasure, 2010 European Conference on Research in Computer Security, ESORICS'10, pp. 643–662.
- [101] V. Gratzner, D. Naccache, Alien vs. quine, *IEEE Security and Privacy* 5 (2007) 26–31.
- [102] J. McCune, A. Perrig, M. Reiter, Seeing-is-believing: using camera phones for human-verifiable authentication, *IJSN* 4 (1/2) (2009) 43–56.
- [103] N. Saxena, et al. Secure device pairing based on a visual channel, *IEEE Symposium on Security and Privacy*, 2006.
- [104] A. Kumar, N. Saxena, G. Tsudik, E. Uzun, Caveat emptor: a comparative study of secure device pairing methods, *IEEE PerCom* (2009) 1–10.
- [105] A. Kobsa, R. Sonawalla, G. Tsudik, E. Uzun, Y. Wang, Serial hook-ups: a comparative usability study of secure device pairing methods, *SOUPS* 2009.

- [106] R. Kainda, I. Flechais, A. Roscoe, Usability and security of out-of-band channels in secure device pairing protocols, SOUPS 2009.
- [107] R. Prasad, N. Saxena, Efficient device pairing using human-comparable synchronized audiovisual patterns, ACNS 2008.
- [108] N. Saxena, M. Uddin, Automated device pairing for asymmetric pairing scenarios, ICICS 2008.
- [109] C. Soriente, G. Tsudik, E. Uzun, BEDA: button-enabled device association, Workshop on Security for Spontaneous Interaction (IWSSI) 2007.
- [110] C. Soriente, G. Tsudik, E. Uzun, HAPADEP: human-assisted pure audio device pairing, Information Security Conference (ISC) 2008.
- [111] R. Nithyanand, N. Saxena, G. Tsudik, E. Uzun, Groupthink: usability of secure group association for wireless devices, ACM UbiComp (2010) 331–340.
- [112] R. Kainda, I. Flechais, A. Roscoe, Two heads are better than one: security and usability of device associations in group scenarios, SOUPS 2010.
- [113] Y. Lin, et al., SPATE: small-group PKI-less authenticated trust establishment, IEEE Transactions on Mobile Computing 9 (12) (2010) 1666–1681.
- [114] O. Chen, et al., GAnGS: gather, authenticate 'n group securely, ACM MOBICOM (2008) 92–103.
- [115] C. Castelluccia, P. Mutaf, Shake them up!: a movement-based pairing protocol for CPU-constrained devices, ACM MobiSys (2005) 51–64.
- [116] Louis Atallah, Guang-Zhong Yang, The use of pervasive sensing for behaviour profiling—a survey, Pervasive and Mobile Computing 5 (5) (2009) 447–464.
- [117] M. Conti, M. Kumar, Report of the NSF/EU Workshop on Future Directions in Pervasive Computing and Social Networking for Emerging Applications, Mannheim, Germany, March 2010. <http://www.percom.org/NSF-EUWorkshop/>.
- [118] C. Bettini, O. Brdiczka, K. Henricksen, J. Indulska, D. Nicklas, A. Ranganathan, D. Riboni, A survey of context modelling and reasoning techniques, Pervasive and Mobile Computing 6 (2) (2010) 161–180.
- [119] N. Roy, G. Tao, S.K. Das, Supporting pervasive computing applications with active context fusion and semantic context delivery, Pervasive and Mobile Computing 6 (1) (2010) 21–42.
- [120] A.C. Santos, J.M.P. Cardoso, D.R. Ferreira, P.C. Diniz, P. Chaaínho, Providing user context for mobile and social networking applications, Pervasive and Mobile Computing 6 (3) (2010) 324–341.
- [121] <http://sensorlab.cs.dartmouth.edu/NSFPervasiveComputingAtScale/> NSF Workshop on Pervasive Computing at Scale, Seattle, Washington, January 2011.