Pervasive Adaptation

The Next Generation Pervasive Computing Research Agenda

A book by *Th. Sc. Community*

Legal notice

Lead editor / coordinator

Univ. Prof. Dr. Alois Ferscha
Institute for Pervasive Computing
Johannes Kepler University Linz
Altenberger Straße 69
A-4040 Linz
Email: ferscha@pervasive.jku.at

Copyright

Acknowledgments

This work was supported under the FP7 ICT FET (Future Enabling Technologies) programme of the European Commission under grant agreement No 213967 (PANORAMA).

ISBN 978-3-200-02270-6

Foreword

In 2011, now 20 years after M. Weiser's "The Computer for the 21st Century" (1991), the vision impacting the evolution of Pervasive Computing is still the claim for an intuitive, unobtrusive and distraction free interaction with omnipresent, technology-rich environments. In an attempt of bringing interaction "back to the real world" after an era of keyboard and screen interaction (Personal computing), computers are being understood as secondary artefacts, embedded and operating in the background, whereas the set of all physical objects present in the environment are understood as the primary artefacts, the "interface". Over it's more than two decades of evolution, the field has been undegoing three generations of research challenges fertilizing Pervasive Computing: The first generation aiming towards autonomic systems and their adaptation was driven by the availability of technology to connect literally everything to everything (Connectedness, 1991-2005). The second generation inherited from the upcoming context recognition and knowledge processing technologies (Awareness, 2000-2007), e.g. context-awareness, self-awareness, resource-awareness, etc. Finally, a third generation, building upon connectedness and awareness, attempts to exploit the (ontological) semantics of Pevasive Computing systems, services and interactions (i.e. giving meaning to situations and actions, and "intelligence" to systems) (Smartness, 2004-). While Pervasive Computing research has its success in the first, partly also in the second generation, the third generation is evolving as we speak.

Preliminarily suffering from a plethora of unspecific, competitive terms like "Ubiquitous Computing", "Calm Computing", "Hidden or Invisible Computing", "Ambient Intelligence", "Sentient Computing", "Post-Personal Computing", "Universal Computing", "Autonomous Computing", "Everyday Computing", etc., the research field has now consolidated and codified it's scientific concerns in technical journals, conferences, workshops and textbooks (e.g. the Journals on Personal and Ubiquitous Computing (Springer Verlag), Pervasive and Mobile Computing (Elsevier), IEEE Pervasive, IEEE Internet Computing, Int. Journal of Pervasive Computing and Communications (Emerald), or the annual conferences PERVASIVE (International Conference on Pervasive Computing), UBICOMP (International Conference on Ubiquitous Computing), MobiHoc (ACM International Symposium on Mobile Ad Hoc Networking and Computing), PerComp (IEEE Conference on Pervasive Computing and Communications), ISWC (International Symposium on Wearable Computing), IWSAC (International Workshop on Smart Appliances and Wearable Computing), MOBIQUITOUS (Conference on Mobile and Ubiquitous Systems), WMCSA (IEEE Workshop on Mobile Computing Systems and Applications), AmI (European Conference on Ambient Intelligence), etc. - with a continued growth in the number of related conferences all over the world. This process of consolidation is by far not settled today, and the focal question that raises after two decades of Pervasive Computing concerns it's future research challenges and roadmap.

Within the FET (Future Emerging Technologies) work programme discussion on new directions for ICTs in FP7 of the European Commission, the question after the next generation Pervasive Computing research challenges has gained momentum. FET acts as a pathfinder for the information and communication technology programme of the EC by fostering novel non-conventional approaches, foundational research and supporting initial developments on long-term research and technological innovation. FET structures research in a number of proactive initiatives, which typically consist of a group of projects funded around a common theme. The themes are shaped through interaction with the research community, and focus on novel approaches, foundational research and initial developments on long-term research and technological innovation.

Specifically, the FP7 FET proactive project PANORAMA (Pervasive Adaptation Network for the Organisation of the Research Agenda and the Management of Activities, FP7 ICT Call-2, FET proactive / Goal 8.3: Pervasive Adaptation) picked up on the challenge of identifying the new trails of Pervasive Computing research. To this end, a structured "Research Roadmap" solicitation process involving some 240 top researchers in the field has been conducted over the past three years. The result of this process is manifested in the Pervasive Adaptation Research Agenda Book (www.perada.eu/research-agenda), which you have in your hands now.

The Pervasive Adaptation Research Agenda Book (Web) Contribution Portal (seamlessly integrated in the community platform www.perada.eu/research-agenda) has been developed as the mechanism to collect and discuss opinions on future Pervasive Computing research challenges with the following key features and services:

- Easy of use: The solicitation process heavily depends on the participation of the community. Therefore the portal was designed to be enjoyable, and to offer functionality at high levels of user convenience.
- Uniform article format: Contributions to the book were to be written by a great number of authors, so it was imperative to enforce strict guidelines. On the portal this is done by providing an input form with distinct fields and input checks. The article body may only be written in plain text (no formatting options, e.g. bold or italic).
- Supporting the review and consolidation process: The portal enables the community not only to author research challenge statements, but also to provide feedback on others. This can be done verbally, by composing a comment, similar to a classic forum, or by voting on the two properties "agreement" and "importance". If several users give a positive vote to the agreement property of a statement, this indicates topicality and attraction, on the other hand, if many negative votes cumulate, a rework might be in order. The importance property is initialized with an assessment by the author (e.g. "impacting" or "pragmatic") and may change due to votes by others.
- Visualise semantic relations: We require authors to specify keywords, an importance assessment and a hint of the time frame the proposed statement would challenge research (e.g. short term, long term). Keywords are clustered by the power of their connections among each other and emphasized in relation to the number of mentions.
- Generate a printable book (PDF): To strengthen the association between the online articles with the printed version of the Research Agenda Book, it is always possible to download a PDF containing all the current articles in a printer-friendly format. A smart versioning system is provided to tell if a book version is outdated.

By June 2011, the Research Agenda Book Contribution Portal has collected about a hundred research challenges articulated by leading experts in the field - on a personal invitation basis, and discussed by the community via the web portal. In the frame of a special session "Pervasive Socio-Technical Fabric" at FET 2011 (The European Future Technologies Conference and Exhibition, 4-6 May 2011 Budapest, Hungary), the Research Agenda Book was presented to a broad audience, and discussed and quality controlled by invited, renowned research pioneers in the field (Norbert Streitz, Coordinator of the EU Disappearing Computer Initiative, Albrecht Schmidt, Figurehead of the European Pervasive Computing Community, and Nigel Davies, Editor in Chief of IEEE Pervasive).

The Research Agenda Book today appears as a reflection of the next generation research challenges in Pervasive Computing, as identified and debated by the scientific community. Therefore, both the digital (web contribution protal), as well as it's physical manifestation (the printed book) appear as authored by Th. Sc. Community. Unique to the Research Agenda Book is it's coexistence in the digital and physical realm, allowing it to continue expanding it's body of contributions also in the future via submissions to the web portal, and generating the corresponding new printed release instantly - with a single "click" from www.perada.eu/research-agenda.

Linz, June 2011 — Alois Ferscha

Alois Ferscha
(Coordinator Pervasive Adaptation Research Agenda Book)

Institute for Pervasive Computing
Johannes Kepler University Linz
Altenberger Strasse 69
A-4040 Linz

Phone: +43-732-2468-8555
Mobile: +43-699-1111-1010
ferscha@pervasive.jku.at
http://www.pervasive.jku.at

Chapter I: Autonomous Adaptation

Chapter II: Adaptive Pervasive Ensembles

Chapter III: Emergence and Evolvability

Chapter IV: Societies of Artefacts

Chapter V: Dependable Pervasive Systems

Chapter VI: Pervasive Trust

Chapter VII: Human-Centric Adaptation

Chapter VIII: Sociotechnical Systems

Chapter IX: Quality of Life

Author Index

Keyword Index

I — Autonomous Adaptation

In pervasive computing we face a rapidly evolving spectrum of an increasing number of services available to anyone, at anyplace, anytime. In everyday life, the average time spent in using such pervasive services ranges from tens of seconds to a few minutes, while they are involved in a growing number of daily, weekly or monthly activities. Clearly, the diversity of platforms, services, tasks and interactions styles introduce serious usability challenges. Moreover, for users involved in such situations, not only speed and accuracy are crucial, but all potential delays accumulate to the rest of people waiting at the same service point. In such sessions, a considerable amount of time is spent by users to repeat sequences of steps for performing the same tasks across independent sessions, such as choosing preferred operations and inputting all required parameters. In this context, adaptive automation is considered a critical element of pervasive adaptation, combining adaptive interaction, adaptive composition, and task automation, by involving knowledge regarding user profile, intentions, point of use, and past use sessions.

It is believed that the ability to record activity traces, recognise patterns of use, identify personal habits, detect repeated tasks, and adaptively apply automations will play a key role towards the broad acceptance and proliferation of pervasive services in the near future. Due to the huge growth of the end-user population, traditional centralised processing and storage approaches regarding user-oriented information may become impractical, even unethical, putting forward the need for personal processing and knowledge resources. While such living, evolving, and personal memory can support advanced future pervasive adaptations, issues such privacy and security become more prominent than ever.

Challenges: privacy prevention, service composition, interface adaptation, task automation, session traces, activity patterns, habits recognition, personalized processing, personal memory.

Adaptive Automations for Pervasive Services

Anthony Savidis

Autonomous Adaptation

Academic Contribution
medium (2) Agreement
long term (>3 years) Duration
challenging (15) Importance

"Greener than Thou" - Pervasive Adaptation: it's Greening Time!

Daniele Miorandi

Autonomous Adaptation

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

impacting (10) Importance

One of the main features of pervasive computing systems is expected to be their ability to seamlessly adapt, in a fully autonomic way, to different operational conditions. So far, this has been mainly intended as the capability to carry out tasks in the presence of changing factors such as, e.g., connectivity among devices, availability of resources (including: computing, storage, energy, atomic service components), requests and expectations from end-users and applications etc. Such an ability can be understood as a way of self-optimizing the system's internal configuration in order to reach a target objective under different environmental conditions.

While self-adaptation is a challenging research theme pre se (How to adapt in a distributed and unsupervised fashion, relying on local interactions only and in the presence of noisy, inaccurate and delayed information?), its use for global optimization has received little attention so far. Further, researchers have mostly focused on using it for enhancing the system's performance according to some application-dependent metrics. At the same time, over the last 24 months we have witnessed an exploding interest in solutions for limiting the power consumption of ICT systems. This trend follows both an economic driver (need to limit the operational expenditures related to energy consumption by operators and/or device owners) as well as a social one (need to limit the environmental impact, in terms of carbon footprint, of ICT equipment). Pervasive computing will not be exempt from such trend. Massive proliferation of computing in wide-area environments may heavily impact the ecological sustainability of innovation in and through ICT.

The question now is: what if we design pervasive computing systems in such a way that they self-optimize at run-time their working parameters and internal configuration in order to minimize their global carbon footprint while satisfying constraints related to the performance achievable by applications running into it? How would such systems look like? What is the impact achievable in terms of reduction of greenhouse gases emission? These questions imply a shift in perspective on how we architect pervasive computing systems. The goal is not any longer to design self-optimizing strategies aimed at maximizing performance, maybe respecting some constraints related to resources usage but, rather, to optimize for minimal energy consumption while respecting some quality indicators on the performance offered to running applications. One interesting remark is that some approaches used in the design of solutions for pervasive adaptive systems, namely the ones based on analogies to and/or inspired by biological systems, appear to be inherently suitable to tackle this new perspective. Biological organisms have always striven for gathering the energy necessary to ensure their survival. More energy-efficient organisms are favored by evolution, as they are able to better survive in environmental situation where energy represents a scarce resource. In other words, many biological processes are already optimized for operating at high level of energy efficiency. Turning them into purposeful design patterns for "self-greening" pervasive computing systems represents one of the grand challenges ahead of us.

Development of Self-Image in Pervasive Systems

Emma Hart

Autonomous Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

PerAda systems are a special case of collective adaptive systems which have particular constraints e.g. they are networked and highly distributed; they involve interaction with humans; they are large scale; the boundaries of systems are fluid; their context is dynamic; and they operate using uncertain information. In order to achieve their ultimate goal of adapting seamlessly to their users and to deliver the expected quality of service at all times, we propose that these systems must exhibit self-awareness. To be self-aware, a system must have an internal self-image: moreover, that image must be dynamic as a PerAda system itself has a dynamic and fluid composition. Borrowing terminology used in both neurology and immunology we propose that collective systems must develop their own homunculus. Stemming from the Latin term "little man", the neurological homunculus refers to the functional map of the body represented in the brain which determines how it interacts with its environment; the immunological homunculus refers to the immune system's internal image of the body. Both the neurological and immunological homunculi can be viewed as a dynamic representation of the body and its various states - the representation is a reduced, virtual image of the system in that it contains far less information than the total amount of information contained in the body itself, but encodes the functional essence of the body's state. A Pervasive Adaptive system must similarly develop its own internal self-image in order that it can both maintain itself and collectively adapt; however, in contrast to the immunological or neurological homunculi, a pervasive system does not have a defined boundary. This fluidity of boundary requires the concept of self-awareness to be completely reviewed when "self" is a fluid entity.

Pervasive Adaptation in Robotic Systems

Eugen Meister

Autonomous Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Pervasive adaptation is a field of research takes its place in a rapidly growing, highly interdisciplinary field, joining contribution from areas as diverse as computer science, engineering, biology, neuroscience, synergetics, cybernetics and many other areas of research. Scalability and adaptability are two big key challenges that emerge in such systems and also cause the most difficulties in the developmental process.

Good solutions for adaptation that are able to handle such complexity are rare and it is a challenge for the future research to develop new methods and techniques for it. To make systems adaptable to some external disturbances is often not enough. We distinguish between different levels of adaptation. On the parameter level we speak about typical adaptation processes, however if the internal/external influences are not predictable or out of any boundaries and hence the whole structure of the system needs to be reorganized, then we need additional structural adaptation level which is currently very challenging to achieve [Levi & Kernbach, 2010].

In the present research, modern robotics overlaps more and more with the research field of pervasive adaptation. In order to achieve adaptation on different levels (structural, functional and behavioral), we need to allow adaptation both in hardware and in software design. Adaptable hardware is one of the most challenging problems in this field. On the level of microprocessors and microcontrollers, there exist only a few solutions how to reconfigure the hardware during the runtime according to the application needs (e.g. programmable embedded system-on-chip - PSoCs [Cypress]). To make mechatronical parts adaptive, it requires also more research and engineering effort. Modular design could be a solution for it, not to forget the research for the new materials (e.g. memory shape alloy, rapid prototyping).

Assuming that the hardware is able to adapt and achieve a high degree of developmental plasticity, the next challenge is to control such systems. New principals for adaptive controllers need to be developed. During the last decades we have recognized that applying neither classical nor bio-inspired methods are sufficient enough. Hybrid methods could help to bridge the gap between these disciplines but are still quite rare. It should be more forwarded in the future.

References:
[Levi & Kernbach, 2010] Levi, P,. & Kernbach, S. (eds). 2010, Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution. Springer-Verlag.
[Cypress] http://www.cypress.com

Cooperative data streams analysis and processing

Giorgio Ausiello

Autonomous Adaptation

Academic Contribution

medium (0) Agreement

mid term (>1 year) Duration

challenging (15) Importance

Future pervasive systems will have to perform tasks in a distributed and autonomous way, possibly in a hostile environment. A typical scenario would contemplate hundreds or thousands tiny processing units with small local memory, that cooperate in order to perform, in an autonomous and adaptive fashion, the planned actions. In order to do so they will have to be capable to perceive signals from the environment through the collection, fusion and analysis of data acquired by a variety of sensors. The problem implied by the distributed processing and fusion of data coming from various sensors is definitely not new. Intelligent systems controlling sophisticated manufacturing plants or fleets of autonomous robots exploring an unknown territory should face similar problems. The novelty that derives from addressing such issues in the context of pervasive autonomous systems is related to two major constraints: memory size and energy consumption. A system as the one depicted in the above scenario may count on a very limited amount of storage that does not allow the saving of large quantities of data with the aim to proceed in a subsequent moment to their analysis and processing. Each unit will have to process data on the fly with a small support of central storage (say of constant or sublinear size) and limited or no access to secondary storage. Decisions on which data have to be stored and which data can be deleted have to be taken in a very quick and irrevocable way, possibly in cooperation with other units. It has to be noted that under such scenario data fusion and analysis is replaced by a different paradigm. Each unit acquires an approximate view of a data stream (sketch) and fusion is then performed among data sketches. On the other side, also the characteristics of distributed computing change in this context since, as it is well known, communications between units is expensive in terms of energy and for this reason communication has to be used in a very conservative way. More generally energy saving requires to limit: (i) number and length of messages exchanged among units, (ii) access to secondary storage (if any) and (iii) use of expensive processing operations. In conclusion data fusion and analysis in the contexts of autonomous adaptive pervasive systems exhibits a variety of problems and represents a very challenging area for theoretical and applied research in the near future.

Future Living Environment

Gordon Cheng

Autonomous Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

Challenges in the realisation of systems that make better use of the pervasive computing, available to us now and in the future shall revolutionise ways humans and machines functions our world, enhancing on our lives in every way. Fundamental needs shall go towards the seamless interfaces that go about in making such technological-life integration possible. The research direction this shall forms the necessity of seamless integration of sensors and actuators in the real world - distributed seamlessly and transparently. Research shall go into the advancements of such technologies that bring about the progress of systems that can richly respond to the real world.

More specially:
1) rapid/compact/cheap sensing in everyday environments;
2) the advancement of processing and management of largely distributed pervasive sensors-actuators networks;
3) the cognitive acceptability of such technology shall need to be validated and examined;
4) with usability studies that account for life-long adaptation across time and space.

Through these technological advancements, that will empower us to interact with the world in a nature and unobtrusive manners, in enhancing our daily abilities. Thus, providing long term benefits to our society as a whole.

Intelligent Pervasive Adaptation in Shared Spaces

Hani Hagras

Autonomous Adaptation

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
challenging (15) Importance

Adaptation is a relationship between a system and its environment where change is provoked to facilitate the survival of the system in the environment. Biological systems exhibit different types of adaptation so as to regulate themselves and change their structure as they interact with the environment. The dynamic and ad-hoc nature of pervasive computing environments means that the environment has to adapt to changing operating conditions and user changing preferences and behaviours and to enable more efficient and effective operation while avoiding any system failure. Thus there is a need to provide autonomous intelligent adaptive techniques which should be able to create models which could be evolved and adapted online in a life learning mode. These models need to be transparent and easy to be ready and interpreted via the normal user to enable the user to better analyze the system and its performance. These intelligent system should allow to control the environment on the user behalf and to his satisfaction to perform given tasks. The intelligent approaches used should have low computational overheads to effectively operate on the embedded hardware platforms present in the everyday environments which have small memory and processor capabilities. In addition, the intelligent approaches should allow for real-time data mining of the user data and create on-the-fly updateable models of the user preferences that could be executed over the pervasive network. Moreover, there is a need to provide an adaptive life-long learning mechanism that will allow the system to adapt to the changing environmental and user preferences over short and long term intervals. There is a need also to provide robust mechanisms that will allow handling the varying and unpredictable conditions associated with the dynamic environment and user preferences.

Most current pervasive systems research have targeted single user pervasive spaces or where the user preference is satisfied in isolation from other occupants. However, this is a major problem is pervasive system where we need first to capture the user preference (which is also changing over time, context and activity) and then we need to formulate models that aggregate the different user preferences to allow a consensus for the environment operation. One area that can help with this challenge is the computational intelligence area which includes fuzzy logic systems, neural networks and evolutionary computation. This challenge need to be addressed inorder to allow the realisation of adaptive pervasive environments.

Homeostatic PerAda Systems

Jon Timmis

Autonomous Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

The biological term Homeostasis refers to an organism's ability to maintain steady states of operation in a massively changing internal and external environment. We propose that PerAda systems should afford the same ability. PerAda systems which are homeostatic offer great potential as despite various failures, user interactions, requirements to adapt and adverse environmental conditions, the system will attempt to continue to operate. A homeostatic system should be able to predict potential problems, and take corrective action; have a combination of pre-programmed and adaptive responses to cope with changing environmental conditions and acclimatise to changes in operational conditions to afford continued system operation.

One possible mechanism for achieving homeostasis is the integration of concepts from neural, immune and endocrine systems. In certain natural systems, such a super system coordinates a variety of complex functions such as: fault tolerance, danger perception, self/non-self discrimination, regeneration, cooperative actuation and perception as well as collective decision-making. In terms of the need for detection, adaptability and controllable long-term autonomy, an integrated immune-neural-endocrine architecture could potentially offer a great deal. Each of the constituent parts could be used not only in isolation, for example, using immune systems for error detection, or neural networks for control, but the integrated whole should afford a level of homeostasis beyond what is currently achievable: however, efficient implementations of such an architecture would be essential to allow for low demands on power resources on the PerAda unit.

Such PerAda systems will be integral to the everyday life of individuals. Therefore, a case will need to be made to give confidence to the validity of the systems developed. It is quite likely that these systems will be too large, or complex, to be validated formally, so alternative approaches may be required. One such approach would be to adopt argumentation techniques developed within the safety critical community. Such approaches make a safety argument for a system, such as the control system for an aircraft. These techniques should be adapted for use in PerAda systems, to construct safety argument structures that give confidence in the operation of the system. In order to PerAda systems to be taken into more critical areas of support, such development is essential.

SensorFly: A Flying Pervasive Sensor Network

Pei Zhang

Autonomous Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

In an emerging class of highly dynamic sensing applications, such as detecting survivors in the aftermath of earthquakes or sensing radiation inside a nuclear power plant, presents new challenges. Human involvement in these situations is often dangerous. The deploying environment changes rapidly, and these applications often require fast discovery. Due to the lack of controlled mobility, previous sensor network platforms are unsuitable for this class of applications.

Robotics platforms have approached these applications with increasingly miniature flying and swarming platforms. However, these platforms have significantly higher costs due to the array of sensors needed for navigation and communication. Because of the fundamental nature of these challenges, there is a need for a new approach to this problem. This approach needs to leverage new mobility technology while meeting stringent constraints associated with low-cost sensors.

To solve these problems, we are developing SensorFly, a first of its kind system based on miniature helicopter-based nodes. The helicopter's total weight of only 30 grams (roughly equivalent to 10 pennies) severely reduces their per-node control, processing, and sensing capabilities. Each node can sense only their orientation, height, and one other input (such as image, sound, light, etc.). Therefore, instead of one robotic node with an array of expensive and accurate sensors on one device, SensorFly divides inexpensive sensors onto multiple networked mobile nodes. Due to the large number of mobile nodes, collaboratively SensorFly can provide faster and more detailed coverage than traditional approaches. In addition, we can fully leverage the nodes' robustness to collisions for lowering navigational sensing needs and for sense obstacles through "touch" (i.e. the nodes senses collision with obstacles). Research for this novel platform is on going in areas of task partitioning, sensor fusion, and coverage. Using networked collaboration and mobility; SensorFly offers a new "cloud sensing" paradigm to collaborative sensing by combining sensor network and robotic platform approaches. In particular, there are several challenges:

— Mobility-aware network setup using shared ambient environmental context such as sound and radio propagation.

— Mobility aware collaborative sensing algorithms

— Sub-meter localization using mobility and ambient context.

— Location-aware, mobility-aware sensing, task distribution and processing

Platforms collaboration in fast deployable systems, like SensorFly, is the future trend in applications where reliability, accessibility and speed are needed. As technology develops and more ambient and ubiquitous applications begin to evolve, the research results that stems from these projects will provide a fundamental basis for future work.

Adaptive Context-aware Heterogeneous Internet-connected Object Networks

Symeon Papavassiliou

Autonomous Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

Among the key aspects of the Internet of Things (IoT) is the integration of heterogeneous objects in a distributed system that performs actions on the physical world based on environmental information gathered by sensors and application-related constraints and requirements. Such an approach is used for the collection, propagation and processing of information and the facilitation of inference and reasoning on the basis of this information. Network design, management and optimisation will be assumed to be performed not just by taking "network-bound" criteria into account, but also considering the nature of information and its intended processing. The mechanisms for this are intended to be generic (i.e., application independent), although of course the reasoning is always application-specific. Such an approach allows to merge two main elements of interconnected object networks, that of information-centric character and node-centric physical world connectivity.

Given such an evolution and taking into account the strong trend of intelligent networked wireless device proliferation, in the future a plethora of enabled devices, sensors and/or objects will act in an autonomic fashion with varying levels of intelligence and capabilities, that will however need to collaborate to accomplish actions and gains that cannot be accomplished without such collaboration. Thus, the corresponding interactions will keep increasing and become more complex, giving rise to a natural trade-off between the gain from collaboration versus the cost of collaboration. As a result the appropriate balance, in a dynamic and operationally efficient way, among the following two basic challenges is required: a) Given behaviour (e.g. application oriented characteristics and requirements), what structure (subject to constraints) gives best performance? b) Given structure (and constraints) how well can this behaviour be executed?

Following this paradigm we consider the future Internet-connected world not simply as distributed collection of asynchronous nodes, but as asynchronous, feedback (many loops), dynamic systems, and distributed asynchronous active databases and knowledge bases. Network Science has strongly emerged as a multi-discipline research area, devoted to the identification, modeling and analysis of such complex networks arising in all kinds of natural and artificial settings.

To cope with this evolution and vision it is required to traverse the whole path of interactions of digital world and physical world following a two way (bidirectional) path: bottom-up, for aggregating and processing data collected by physical sensing devices and objects using it for situation awareness and decision making at the higher level; top-down, for compiling application demands into sensing requests and control actions over both the virtual and/or physical infrastructure through topology modifications, taking into account arising properties and features from social networks in order to improve performance.

II — Adaptive Pervasive Ensembles

We propose the novel computing and analysis paradigm of Dual Spatial Reality. Developing and improving formalizations and modeling frameworks for real world processes to save resources of various kind has always been one of the major scientific goals of computer sciences. In the last decade these findings have led to the conceptualization of systems modeling certain aspects of the world, such as the information and production chain in a company (->ERP Systems) or the representation of traffic flow in a city (-> traffic management systems). These systems are based on abstractions of reality and often use highly simplified assumptions. They can only be operated on and interpreted by domain experts, which are familiar with the simplifications and with the domain itself. The systems are typically not connected - yet, it is clear that they could benefit from each other. For example, the traffic management system could benefit from a system which models the local climate to help forecast road conditions or the degree of pollution in certain areas of the city. Various active research communities are developing concepts and theories to overcome these system boundaries, including research on semantic modeling and integration, sensing and interpretation, and system architecture. At the same time, the ongoing miniaturization of computing components and the development of sensors and actuators allows more and more computing hardware to be embedded directly into the real world, which in turn allows to measure and to react to conditions of the real world in real-time. This connection is highly beneficial: Problem states can more easily be transferred from the real into the digital domain (by sensors) and the results of reasoning processes of the digital domain can directly be transferred back into the real world (by actuators). This tight connection between the digital and real world is what will lead to a Dual Spatial Reality. The main goal of this research proposal is to work toward this vision by improving

— information extraction from the real work through sensors (-> pattern recognition and sensing);

— geospatial process models and integration of information (-> semantic integration);

— models of economic processes and of the flow of goods and objects (-> information sciences and logistics);

— system architectures and programming languages that help to integrate

Computational Methods for Dual Spatial Reality to support Urban Every Day Tasks

Antonio Krüger

Adaptive Pervasive Ensembles

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

impacting (10) Importance

different existing systems (-> research on SOA and software engineering);
— interaction with information in both realities (-> computer graphics, tangible user interfaces, novel interaction paradigms).

To exemplify and test our ideas, one idea could be to develop a Dual Spatial Reality on an urban scale to show how quality of life is improved by a tight and two-way connection of the real city and its digital counterpart. We will investigate these benefits at multiple scales, ranging from transportation in the city to efficient smart factories and innovative forms of medical diagnosis and treatment. The idea has a strong potential to be expanded toward the natural and social sciences, for example including research on novel sensors and investigations of societal impacts of this novel computing paradigm.

Pervasive Vehicular Environments

Djamshid Tavangarian and Robil Daher

Adaptive Pervasive Ensembles

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

challenging (15) Importance

Vehicular Communication Networks (VCNs) have emerged as a key technology for next-generation wireless networking and become the basic communication platform for Intelligent Transportation Systems (ITS) services. VCNs represent the main infrastructure for providing a Pervasive Vehicular Environment (PVE), which brings together information and telecommunication technologies in a collaborative scheme of applications to not only improve the safety, efficiency and productivity of transport networks, but also to provide and enhance multimedia and data services for both mobile users and vehicles. However, VCN characteristics cause several new issues and challenges for operation of pervasive environments, especially when vehicles travel at high speeds of up to 200 km/h and produce thereby a high dynamic network topology. These issues and related challenges, addressed in regard to roadside networks and Vehicular Ad hoc (unplanned) Networks (VANETs), including Vehicle-to-Vehicle (V2V) and Vehicle-to-Roadside (V2R) communications, form a promising research field not only for strengthening the long-term evolution of PVE, but also for effectively understanding, among others, the cooperative intelligence should be integrated into future pervasive/mobile environments.

In general, the high dynamic network topology of VCNs represents the main challenge for providing efficient network performance, guaranteeing real-time capability, and offering a high-level security in PVEs. Firstly, although the network performance of current VCN-technologies provides a sufficient platform for ITS services, the long-term adoption of PVEs reflects the necessity for more bandwidth and network availability. Therefore, future PVEs require an interactive heterogeneous network environment including several types of radio access technologies such as DSRC, WLAN, WiMAX and UMTS in order to improve pervasiveness and network redundancy. Secondly, the real-time capability of VCNs represents the main challenge for users' acceptance of PVEs. Due to the lack of research in the field of VCN-specific QoS solutions, several topics, such as message/packet prioritization and dissemination, mobility-aware resource management, routing and seamless handoff are still open to be addressed. Finally, the security and privacy in PVE form very hot topics for future research and are expected to determine the time and scale of PVE integration into current and future network infrastructures. The typical challenge for PVE is to bring together high-level security/privacy and high mobility with minimal degradation of network performance and real-time capability.

Furthermore, adopting PVE without considering the roadside backbone network cannot offer a reliable guarantee for service and access pervasiveness in the whole VCN. Therefore, challenges for designing reliable, modular and real-time capable roadside backbone infrastructures should lead the research topics in the future. Finally, the development of PVE-oriented service architectures forms an essential step toward revolutionizing the service and network design in the future.

Cooperation in Pervasive Computing Environments

Gabriele Kotsis

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

Being able to accomplish tasks in a joint effort was and is a driving motivation for humankind. Thus, from the very beginning, information and communication technology was designed and used to support humans in their cooperative efforts. With the era of pervasive computing, cooperation no longer involves human beings but also software and hardware artefacts, sometimes summarized in the term "machines". We learned a lot from cognitive and social sciences about human interaction which has significantly influenced the way in which we nowadays design and develop ICT supporting cooperation among humans. When machines come into play, a new way of understanding is needed in analysing on the one hand the different aspects of cooperation among machines and humans (H2H/H2M/M2M) and in translating these findings on the other hand in guidelines, principles and methods in system design, development, and evaluation. The variety and dynamics of "machines" present in a pervasive computing ensemble, their different requirements and capabilities along with the different application and user specific needs indicate the need for flexible, extensible, and adaptible cooperation models and mechanisms which current technologies do not offer.

Probabilistic Modeling and Adaptation

Gaetano Borriello

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

As we improve our ability to capture data from myriad sensors and devices, it is now clear that modeling of a users' context and the activities that they are engaged in is a central problem in making more responsive and robust pervasive computing ensembles. Context-awareness is woefully limited in our computing devices and they rarely do the "right thing" or what we would prefer. We need to be able to teach them how we want them to work for us. To date, we have focused on simple sub-problems such as a user's location, or whether they are running/walking/biking, or are in a noisy environment, but these are not context or even the full picture of complex and high-level activities. What we need are the ability to tell if a user is interruptable, what information they likely to need next, what work/play they might be engaged in, and who might be engaged in it with them (just to name a few). What we need is a personal "Radar O'Reilly" - the ubiquitous clerk character of the book, movie, and television series M*A*S*H. Our approaches to date leave much to be desired. They have focused on the deterministic and mostly the one-size-fits-all model. But our sensing is noisy and incomplete and likely to remain so in people's real lives outside of laboratories. Moreover, our models will never cover all possible activities in which humans may engage. A huge challenge for our community is the modeling of this probabilistic information and the uncertainty it implies. But modeling is not the final goal. This will then need to be followed by the development of application architectures that deal with the probabilities directly and, very importantly, adapt to user reactions with as little effort as possible from the user.

Pervasive distributed processing capacity of complex networks

Gianluigi Ferrari

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

The scientific community has been looking with a very strong interest to the fields of ad hoc and sensor networking, from the most various viewpoints: from a networking perspective, from a communication perspective, from a system perspective, from a control theoretic perspective, etc. Likewise, many approaches have been proposed to study these networks, such as: single layer-oriented, cross layer-oriented, graph theoretic-based, bio-inspired, etc. This amazing attention is mainly due to the intrinsic complexity of these networks, which, therefore, opens a myriad of problems to be solved. On the other hand, distributed pervasive processing is the "soul" of all proposed approaches. However, to date a variety of performance results have been presented, but it is not clear "how good" these results are. In other words, the following question cannot be answered: is the performance of a complex network close to its ultimate limits? In order to answer this question, a unifying theory of pervasive distributed processing needs to be developed.

Pervasive distributed processing involves three main phases.
1) Sensing: in this phase, "information" is extracted from the real world, typically by means of sensors. Processing, in this phase, is local: proper operations can be locally carried out on the extracted information (e.g., quantization).
2) Communication/networking: the extracted information may flow across the network, to either reach a specific destination point (a sink) or a specific node/group of nodes. In this case, processing may be carried out "on the way" (e.g., distributed multi-level fusion) or at the sink.
3) Acting: the processed information may be conveyed to the network or part of it, in order to be properly act or react (e.g., efficient information broadcasting, in vehicular ad-hoc networks, to warn nearby drivers of an accident).

On the basis of the above simplified view, one could determine the ultimate (achievable) performance limits, guaranteed by the use of pervasive distributed processing in a complex network, by developing a unifying theory. For instance, the "pervasive distributed processing capacity" (PDPC) of a complex network should be defined and, then, evaluated. The PDPC should consicely reprent the insintric potential of a complex network. Upon clear definition of the PDPC, any feasible distributed processing strategy could be evaluated.

Obviously, the definition itself of the PDPC is very challenging, since it should take into account many aspects: information flow, energy consumption, delay of reaction, resource sharing, etc. However, the rewards, in terms of clear performance benchmarks in the design of effective distributed processing systems, could be tremendous.

H-W2 Motions in Computing

Justin Y. Shi

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
pragmatic (5) Importance

Technology innovations are mixed with three inter-twined motions: How-to, What-to and When-to. While the networked computing apps are catching our imaginations, recognition of these motions of innovation may serve us well while exploring the unknown worlds. As we thrive to solve world's problems using technologies, how to avoid unintended consequences exemplifies the difficulties of technology. Mark Weiser's vision projects a perfectly balanced H-W2 motion in time, where sophisticated computing powers are efficiently delivered with the ultimate convenience, the commercialization is done in the right time that everyone is ready to accept it, and the technology impact is perfectly interfaced with higher level infrastructures that unintended consequences would never occur or at least for a very long time.

Historically, theoretical analysis often preceded practical implementations. Modern technologies have made implementations easier tantalizing ventures. While we are having fun pushing the edges of technologies, the need for theoretical foundations cannot be over-estimated. Although difficult, historically, they have saved us from resource wastes and unintended consequences.

Theories also have limitations. One example is perhaps the LZW compression algorithm. Although Hoffman code is the theoretically optimal, the Lempel-Ziv-Welsh compression algorithm has been shown pervasively useful even the proof is still been sought. The near future How-to challenge is perhaps sustainability: how can we design apps that can scale in performance, reliability and security at the same time? While there is no lack of What-to proposals, we quickly find that all must interact with higher level social, economic and political infrastructures in some way. Finding the unintended consequences is both practically and theoretically challenging, since the art of W's makes billionaires and bankrupts research labs.

A unique feature of pervasive computing is its cross-discipline nature. It provides a fertile ground for theoretical and practical challenges. Keeping a balanced exploration mode is perhaps critical to our future successes. For example, we would be all better off if we knew what makes a viral video viral, what mechanisms could prevent extreme activities, what if all 911 callers are busy, how to monitor people without intruding privacy, and why texting are less susceptible to spams?

Adaptation without Anticipation

Kay Römer and Friedemann Mattern

Adaptive Pervasive Ensembles

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

challenging (15) Importance

A pervasive system should be able to perform a desired function despite changes of environmental conditions or the system state. For example, a distributed visual object tracking system should provide a certain user-defined accuracy even in the case of varying environmental lighting conditions or when some sensors fail. Such pervasive adaptation is typically a distributed problem: adapting the behavior of the system as a whole requires coordinated actions of its distributed components.

Current practice in supporting pervasive adaptation requires the developer to manually specify the adaptation rules of the distributed components. However, the manual translation of desired global behavior into local adaptation strategies is not only intricate, but also requires the anticipation of all conditions and states that could be encountered.

We envision a system design methodology that relieves the developer from these difficulties and potential error sources. In particular, one should be able to specify the desired global system behavior using appropriate high-level specification languages. The pervasive system should then be endowed with an infrastructure to develop adaptation strategies for its components such that the desired global behavior is delivered across all possible situations.

Realizing this vision presents significant challenges at different levels. Although so-called macro-programming systems [1] allow a direct specification of global system behavior, the semantic gap between the user mindset and those languages is still very large. Also, these languages often represent a binary specification of the desired behavior, while many optimization strategies that would be part of the adaptation infrastructure require some form of objective function as a quantitative measure of closeness to the desired behavior.

While optimization techniques such as genetic programming have demonstrated to enable automatic derivation of strategies that meet certain high-level goals [2,3], they are computationally expensive even for simple problems and hence need to be performed offline on high-performance computers. Offline methods, however, need explicit models of the environment and the system which already anticipate all relevant conditions. An alternative are online approaches [4,5] where for example the pervasive componentes themselves would modify initial strategies to match the desired global system behavior using the actual environmental and system states they encounter. In that way, the system should eventually be able to adapt to unforseen states.

While initial steps have been undertaken in support of this vision [1-5], they are little more than proof-of-concepts using much simplified problems.

Significant effort is required to turn this vision into practical solutions.

References:
[1] Luca Mottola, Gian Pietro Picco: Programming Wireless Sensor Networks: Fundamental Concepts and State of the Art, ACM Computing Surveys, to appear.
[2] Arlindo Silva, Ana Neves, Ernesto Costa: Evolving Controllers for Autonomous Agents Using Genetically Programmed Networks, EuroGP 1999.
[3] Thomas Weise, Kurt Geihs: DGPF - An Adaptable Framework for Distributed Multi-Objective Search Algorithms Applied to the Genetic Programming of Sensor Networks, BIOMA 2006.
[4] Philip Valencia: In situ Genetic Programming for Wireless Sensor Networks, Sensys 2007 Doctoral Colloquium.
[5] Suranga Hettiarachchi, William M. Spears, Derek Green, Wesley Kerr: Distributed Agent Evolution with Dynamic Adaptation to Local Unexpected Scenarios, WRAC 2005.

Collective Adaptation in Open Environments - The Challenge of Stability and Conflict Handling

Klaus Herrmann

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

One of the most difficult and at the same time most important goals in designing an adaptive system is stability. Since an adaptive system changes itself (e.g. its parameters and behavior) in response to dynamic changes in the environment, there is a danger that it enters oscillations or a long sequence of successive adaptations without converging to a stable state that fits the requirements. There is a solid theoretic basis for ensuring and proving stability in control theory, and there are tailor-made means proposed in individual adaptive computing systems. However, most of these systems are not designed with openness in mind, i.e. they assume a priori knowledge of all possible effects that influence the adaptation process. But what happens if two or more adaptive systems interact directly or indirectly in the same environment, without having explicit knowledge about this interaction? Pervasive computing is one example where this may happen frequently: different autonomous adaptive systems that have not been designed to know about each other meet (by the means of mobility) and act on the environment, each according to its own plan. E.g. they may allocate specific resources and use services provided in the environment. This has at least an indirect effect on the other systems that do similar things, e.g. because allocated resources are no longer available for them.

Approaches
The problem of resolving such conflicts and avoiding instabilities can be

generally solved in two ways: 1. by offline model-checking, and 2. by online coordination. Offline model-checking tests the overall system in the design phase. This is not viable since the involved systems are not known in advance, which leaves us with some online mechanism that coordinates all interacting adaptive systems as they run. However, such coordination is problematic as it has to solve a very hard problem in a distributed fashion in a dynamically changing environment.

Challenges

A major challenge for the European research on adaptive systems in the next years is to create concepts for enabling autonomous adaptive systems to interact in open environments without ending in conflict or entering instable behavior while individual systems enter or leave the collection of adaptive systems. Adaptivity will only be applicable on a large scale beyond closed laboratory environments if this problem is resolved. This is true for all areas of distributed computing and especially for pervasive systems that inherently operate in an open world.

Interaction in Limbo: dealing with highly dynamic environments and uncertainty

Kris Luyten, Karin Coninx, Joëlle Coutaz and Alexandre Demeure

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
pragmatic (5) Importance

In the ubiquitous computing paradigm the computing paradigm shifted from mostly desktop computing to embedded distributed computing integrated in the user's physical environment. The distributed computing ensembles need to be connected in an ensemble so they act as a whole to support the end-user's activities. This makes the design of interactive pervasive system painfully complex: in contrast with desktop computing systems there are no well-defined goals, tasks and plans. One cannot design for something that can not be defined. User interfaces for pervasive systems are in limbo: they are complete but not finished, they are often useful but not always sound.

For the creation of future pervasive applications, the design should no longer be definitive, but as a starting point for further sculpting and configuration by the end-user. This applies to every step in the creation process of a pervasive system. The boundaries between developer, designer and end-user are fading. End-user drive composition and configuration of ensembles thus having their say in the structure and behavior of the ensemble. There is a parallel one could draw with the DIY (Do-It-Yourself) culture that has recently gained a lot of interest in the community. We need new tools for creating pervasive applications that transcend the typical responsibilities of these formerly in time and space isolated roles (developer, designer and end-user) and make them accessible for the other roles. The importance of the social network is directly related with this evolution: it is reasonable to think that the next generation of end-users, often digital natives but non computer experts, will co-develop new services and pervasive applications.

The cornerstone to allow end-users to have control over a pervasive application are meta-user interfaces. Coutaz defines meta-user interfaces as "a special kind of end-user development environment whose set of functions is necessary and sufficient to control and evaluate the state of an interactive ambient space". We argue a meta-user interface should provide access for querying and manipulating ensembles. Therefor, in an ensemble each component contributing to pervasive applications should be self-described and provided with computer processable semantics. In many ways an ensemble for supporting user activities in a pervasive computing systems is like a mash-up of web services as is common in web 2.0. The semantics provided by these components should be easy to aggregate and query to provide us with the information required for generating the meta-user interface.

Meta-User Interface should expose some specific capabilities in order to be useful for the end-users. They should be intelligible and make the end-user aware of the behavior of the system, even help the user if necessary (e.g. by answering "why did that happen" or "how can I accomplish X" questions). They need a way of visualizing themselves on arbitrary surfaces and devices in the user's vicinity. They need to support both individual users and groups of users. Finally, they should define operations with which user's can manipulate the application specific user interfaces and behavior (migrate, aggregate, share, split, transform,...).

Trusted Short Range Wireless Communication in Pervasive Interaction

Mark T. Smith

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

The wireless landscape continues to change and shape the kinds of interaction models used in adaptive pervasive systems. One of the more complex areas where this is taking place is in the role of short range and near field wireless technologies. As a research agenda in pervasive systems, short range wireless technologies carry the ability to change both human to machine interaction and machine to machine interaction. They have this ability because while being low cost and easily obtained, short range wireless systems are becoming platforms for more than wire replacements for low bit rate data exchange. They represent new human interaction models based on proximity and touch. Possibly more significant is that this proximity and touch model of connectivity can occur between any number of non-human objects and can drive a wide range of actions. Short range wireless systems in the form of RFID or Near Field Communications are also incorporating models of data standardization, security, trust, and authentication, all of which raise interesting design and ethical questions in future adaptive pervasive systems. The core of this research agenda items is in how emerging short range wireless systems will be deployed, and the accountability models needed for their use. For example, with respect to deployment, some short range wireless components are simple enough that it may be possible to incorporate them directly into everyday materials. Another example is concerned with who will control the use of new short range wireless systems. Because the operation of such systems may intentionally occur without a user's awareness, these systems will need to incorporate methods to account for system and user authorization, and to guarantee authenticity of the components that make up any transactions. To account for this, next generation short range wireless systems will employ hardware for trusted or secured data transactions. This represents significant research problems in how secure data elements are designed and used, and where they and the transactions they regulate are located. Future pervasive systems will need to balance users' needs for functionality and trust with commercial requirements for profitability and control, so it is important that research in these systems address how the secure data elements are designed, used, managed, controlled and provisioned. The number of stakeholders in these new systems is potentially large, and so research in these systems will, in addition to wireless and secure element technologies, take into account aspects of economics, business and politics.

Research challenges in pervasive computing

Michel Riguidel

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

Pervasive computing is a smart space concept; it is a spatiotemporal medium of communication, in which a digital ecosystem with multiple granularities, and interactively people and organizations are interdependently evolving and changing, including also natural elements that surround them.

A pervasive system is a digital urbanization, which consists of infrastructures equipped with specific entities that are part of the network and provide mediation services to assist communication (bandwidth broker, location manager, topology nodes memorization) or to facilitate the operation of terminals (configuration, service discovery).

Pervasive computing is not a homogeneous world of free robots to help us in our painful tasks. It is a "situated computer science", a computation, which depends on the function of the immediate environment, the context, the cybersphere size. A pervasive system is made of sensor networks, smart devices, and swarms of chips with scarce resources clinging to the building walls, physical infrastructures which supports calculation, storage and communications, cooperating with the communication functions to facilitate delivery and are used to configure applications according to their geo-located usage. A pervasive system can use external services provided by alternative infrastructures (time, space, radio resource management). Pervasive systems for an individual in an apartment or on a campus are distinct. They adapt to the environment, services and claimed usage. The semantics extends and varies according the niche granularity of the digital ecosystem. The substitution of a common language with different adjusted vernacular dialects, if it stimulates the adequacy of the situation raises insurmountable obstacles. The "smart dust" that is spreading in the urban landscape pose serious security questions.

Pervasive computing can be designed at various scales:
— the planet to observe, analyze, detect, identify, alert, decide, act;
— human communities to collect data in real time;
— a city or neighborhood to help citizens;
— a transport infrastructure to facilitate the travel of passengers;
— a home for education, information, control home automation, daily life;
— a person to monitor health through control mechanisms of digital prostheses in the human body with interactions with the outside, whether the patient himself, a nursing aide at home or a physician connected by telecommunication .

Pervasive computing is a multidisciplinary research in computer science which is transversal to microelectronics, software, but also human and social sciences, legislation since the behavior of actors plays a crucial role in the evolution of this digital reign for manufactured goods, smart walls, agoras (railways stations, airports) and our vehicles (trains, cars).

Beyond the slogans and mercantile jargon, we have not yet dare to break and go to the end of the consequences that pervasivity embrace in reality, especially in models of communication, computation and storage and interaction with physical reality (robots), living reality (computer science for organs, cells, bioinformatics), and invisible reality (nanotechnology).

Challenges of pervasive computing are:
— mobility (with persistent and continuous communication), moving (with intermittent connection) of communicating entities, contextualization and environment learning (service discovery, adaptation, cognitive networking) of ubiquitous computer, remote presence of individuals or presence of fragmented bodies;
— scalability issues for identity management and identity of massive objects,
— fundamental properties of communicating objects: safety, transparency, non-intrusion, non-addiction vis-à-vis the user;
— multimodal interfaces, heterogeneous interactions and negotiations within the assemblage of digital urbanization and the overall composition of different aggregates in the architecture of various systems;
— integration, spontaneous or opportunistic cooperation or collusion and learning capacity at all granularities of all the intricate niches and governance of this complexity;
— cross-cutting aspects of usage taking into account multidisciplinary.

The non-technology aspects cover a wide spectrum:
— Friendliness: Intuitive (emotional) interaction, more general than ergonomics;
— Psychological: discharge phenomena vis-à-vis devices to aid the elderly;
— Legal: safety of robots, intellectual property, the right of access to a digital entity, the right to archive and to forget all the recorded tracks by these objects, the regulation on physical objects;
— Ethics: digital sovereignty, respect for individual privacy, digital dignity with infrastructure for surveillance;
— Policy: freedom of expression for citizens in semi-private spaces.

Fundamentals of Pervasive Nanoscale Communications and Nanonetworks

Ozgur Baris Akan

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

"There's Plenty of Room at the Bottom" was stated by 1965 Nobel laureate physicist Richard Feynman in December 1959 expressing the vision on the possibility of miniaturization of the existing devices to molecular and even single atom level. Our aim is to extend this vision to reveal the answers to the question: is there a plenty of room for a communication network at the bottom? Recent progress in nanotechnologies have yielded the nanomachines, which we define as the particle-level functional units comprised of an arranged set of atoms or molecules that are able to perform the most basic computation, detection, actuation, and communication tasks. Clearly, a single nanomachine has significantly limited processing, and mobility capabilities. Hence, "nanonetworks" are defined as a set of nanomachines communicating with each other to realize a complex objective.

Nanonetworks bring a set of unique challenges due to physical constraints and operational principles of nanomachines. Nanoscale devices also naturally communicate with each other through various mechanisms governed by the laws of quantum physics and electromagnetics. Thus, all known results, traditionally safe assumptions, and fundamental functionalities for classical communication systems are inapplicable to this domain. In this emerging and very interdisciplinary research thrust, we will discover fundamental and practical nanoscale communication paradigms, theoretically model and analyze, and develop frontier communication mechanisms for the realization of nanonetworks. Main nanoscale communication paradigms: - molecular communication, - carbon nanotube (CNT)-based nanoscale-electromagnetic (nano-EM) communication, - nanoscale quantum communication should be investigated. We must explore the molecular communication mechanism, from radically different and pervasive communication perspectives, i.e., diffusion-based communication, molecular neuro-spike communication, Foerster Resonance Energy Transfer (FRET)-based molecular communication.

For truly pervasive networking and practical realization of ultra large scale Internet of Things (IoT) vision, novel techniques to exploit tiny CNT-based nanotransceivers towards realizing nanonetworks must be investigated. We will develop frontier communication techniques for information encoding, modulation, packetization, medium access control, routing and reliable delivery for both static and mobile ad hoc nanonetworks. A new nanonetwork simulator framework, called as NanoNS, will also be developed to validate the nanoscale channel and network models, and evaluate the performance of the communication techniques developed. Interdisciplinary experimental studies will be carried out to validate the developed communication models and techniques in a physical testbed. Nanoscale communication models and techniques that will be developed based on laws of physics, and hence, will be among the first realistic solutions in this very young and interdisciplinary research field.

From Context Awareness to Socially Interactive Computing

Paul Lukowicz, Dirk Helbing and Steven Bishop

Adaptive Pervasive Ensembles

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

The core of current work on context-aware systems involves small groups of devices interacting with a single user or a small group of users in system configurations specifically designed for often narrowly defined applications. On the other hand, smart phones, home automation devices, robots and other intelligent systems are becoming ubiquitous and are increasingly equipped with flexible networking capabilities. Thus, looking at the future of context aware systems, we need to consider large collectives of such devices dynamically forming, cooperating, and interacting with large user populations over a broad range of spatial and temporal scales. Such interactions imply four core research challenges:
1) Opportunistic information collection. Systems need to be able to function in complex, dynamic environments where they have to deal with unpredictable changes in available infrastructures and learn to cooperate with other systems and human beings in complex self-organized ensembles.
2) Collaborative Reasoning and Emergent Effects. Reasoning methods and system models are needed that combine machine learning methods with complexity theory to account for global emergent effects resulting from feedback loops between collaborative, interconnected devices and their users.
3) Social Awareness. Whereas today's context-aware systems are able to make sense of the activity of single users and their immediate environment, future systems should be able to analyse, understand and predict complex social phenomena on a broad range of spatial and temporal scales. Examples of the derived information could be: shifts in collective opinions and social attitudes, changes in consumer behaviour, the emergence of tensions in communities, demographics, migration, mobility patterns, or health trends. One can imagine such systems to form a virtual "social nervous system".
4) Social Adaptation. Beyond adapting to the needs of single users, ICT systems will have to develop the ability to consider the needs of entire communities and social structures. They will have to be able to react to unforeseen events in a way that will have a favorable effect on social processes and phenomena. The adaptation could range from reshuffling of resources (e.g. information sources, bandwidth, distributed computing resources) to enable a better monitoring and management of emerging crisis situations, over the mediation of interactions in and between communities, up to emergency "slow down and ask human" mechanisms, preventing the system from accelerating crises and systemic failures. They will be directed by high-level, human-formulated goals and implemented by bottom-up,

self-organized processes leveraging the systems' social awareness and ability to model complex social phenomena.

In summary, in future we will have to take a more integrative view of ICT as (1) a complex, dynamic, globe-spanning system, (2) composed of billions of entities interacting over multiple spatial and temporal scales, (3) interweaved with society in a multitude of ways across different functional and structural layers. A core question that future research needs to address is how to turn the relationship between the complex, globe spanning ICT infrastructure and human society into one that benefits humanity and leads to more stable, trustworthy, reliable, and inclusive ICT.

III — Emergence and Evolvability

Evolving Artificial Regulatory Networks (EARN)

Ernesto Costa

Emergence and Evolvability

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
challenging (15) Importance

Complexity in biologic systems is a product of evolution, with evolution itself a many dimensions process (e.g., genetic, epigenetic), and of development. If we look inside those systems, what we see is a decentralized, heterogeneous network of interacting elements, capable to respond and adapt to new, unpredictable, and dynamic environments. For those systems to keep working properly the existence of feedback, regulatory loops is crucial. Regulation is the key for homeostasis, and regulatory networks are the expression of complexification in natural systems. There is an increasing interest in bio-inspired artificial regulatory networks with two different, yet complementary, goals in mind: as computational models, leading to a better understanding of the underlying biological mechanisms, and as problem solvers models, used mainly in optimization contexts. Less attention has being paid to the process that led those networks to come to existence, and how it can be translated and used computationally.

Pervasive Computing is facing novel challenges. These are not only technological, i.e., related with computing devices, communications and user interfaces, or human and social, i.e., involving privacy and safety. Some of the issues are architectural and related with foundational concerns. Today, pervasive systems have to be designed, implying the capacity to foresee lots of possible working situations. Systems' intelligence and adaptability is still very limited. Moreover, as the number of components increases, it becomes difficult, if not impossible, to harness the emergent complexity of the whole system. It is natural that we ask ourselves if it is possible to start with a minimal architecture and let the system grow and develop by it- self, as an answer to the environment demands and system's goals. Can such a system go through different phases of maturation? How to achieve progressive self- modification?

At the architectural level, and from an abstract point of view, pervasive systems face the same problems as those that some biological systems have to deal with. They are both a distributed society of networked elements, each of them with some communication, processing and actuation capabilities. If nature was able to find simple solutions for complex problems, can an artificial system do the same? Understanding how regulatory networks developed in nature, how they find specific solutions from the same basic building blocks, is a first step towards answering that question. From that we can proceed to build a computational model of this process, and apply it to the construction of pervasive systems. This is a demanding endeavor, both scientific and technological, whose success is full of consequences.

Emerging design

Giuseppe Persiano

Emergence and Evolvability

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
challenging (15) Importance

Examples of complex systems with unpredictable yet robust behavior are abundant in Nature and, unlike human-made artificial systems, do not seem to be the fruit of design. In such systems, individual tiny components simply respond to local stimuli and eventually position themselves so to maximize their benefit. The benefit can be of economic nature (with higher benefit corresponding to a larger share of some basic resource) as is the case of human societies or of different nature (for example some complex physical systems can be seen as made of tiny particles each trying to minimize the potential energy). The apparent lack of design is instead hiding the simplest possible design: attach a benefit to each component and let the system evolve to a stable configuration. The analysis and the prediction of behavior of such a complex system then boils down to the study of its set of stable configurations. Similarly, to design a system that has a desired behavior one simply designs a system for which the desired behavior is a stable configuration.

It is our thesis that future complex systems will not be designed following a top-down approach but rather by "instructing" each tiny component from a network of components to increase a specific benefit. The tiny component will dynamically adapt to the changing environment so to increase their benefit and this adaptive rules will be of a local nature (that is depend only on the neighborhood of the component) and will not impose a large computational load on the component. The desired behaviour will then *emerge* from the system. A primary example of this approach is how the Web has evolved: no central authority has designed how web pages should link to each other. Rather, a multitude of agents have created web pages and added links so to maximize the usefulness of the pages and the likelihood that each page will be read. As the Web changes (because new pages are created or old pages become obsolete), so do links from any given page.

For such a new design approach to complex systems be viable we need to understand

1. given a set of simple local rules, what global stable configuration the

system will reach;
2. how much time it takes for the system to reach a stable configuration.

We stress that classical approaches to Economic and Game Theory do employ the concept of stable configuration (one of the primary examples being the concept of Nash equilibrium) and they concentrated on the study of the system once the equilibrium has been reached. Little attention has been devoted on how (and how fast) such an equilibrium can be reached. We expect the full development of this approach will use concepts from Statistical Mechanics, Learning Theory and the Theory of Markov Chains.

The Evolution of Images

Oliver Bimber

Emergence and Evolvability

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
impacting (10) Importance

As many other application-oriented disciplines, pervasive computing relies on the advances of enabling technologies that are made possible by fundamental research fields, such as material sciences, mechanics, physics, optics, electronics, electromechanics, and others. Embedding microprocessors in everyday objects would be unimaginable without the dramatic miniaturization of semiconductor technology within the previous decades. Today, electronic circuit assemblies (e.g., based on carbon nanotubes) can be transparent and flexible, can be printed on polymers, and can be interfaced to nerve cells -- offering as yet unimaginable applications. The potential of future pervasive systems clearly depends on the evolution of such technologies.

Today, several pervasive systems rely on imaging and display technologies. Since the invention of first camera- and display systems, images always have been two-dimensional matrixes of pixels. Even though 3D scene points emit varying light rays in different directions, the lens and the sensor of cameras integrate them to a single pixel. By doing this for all imaged scene points, we end up with nothing else than a 2D image -- having lost most of the scene information. The same applies to many displays -- although in a dual way. Pixels of raster-displayed images emit (more or less) the same amount of light in all directions -- giving us nothing more than a 2D image.

What if the notion of images would change once and forever? What if instead of capturing, storing, processing and displaying only a single color per pixel, each pixel would consist of individual colors for each emitting direction? Images would no longer be two-dimensional matrices but four-dimensional ones (storing spatial information in two dimensions, and angular information in the other two dimensions). In electrical engineering, this four-dimensional structure is called photic field. In computer science, it is called light field.

Light fields are radically changing everything that we relate to images -- from photography, over displays to image processing and analysis. But what does this mean for pervasive systems? Embedded light field displays, for instance, will no longer be 2D but will deliver 3D images to an arbitrary number of observers with no demands for glasses or tracking. Embedded light field cameras, as another example, will no longer deliver 2D images, but 4D information that takes scene analysis far beyond of what is possible with any 3D range scanning or 2D imaging technology today.

While first light fields display prototypes have already been introduced in scientific communities and first light field cameras are becoming already commercially available, the real challenge is the processing of light fields: common digital images store mega-bytes of data -- corresponding light fields will require gigabytes.

Since pervasive computing is strongly technology driven, it is worth considering the evolution of enabling technologies to estimate the future impact of pervasive systems in certain application domains. How do pervasive systems advance when images finally evolve to light fields?

Challenges of Pervasive Adaptation: Viewpoint from Robotics

Serge Kernbach

Emergence and Evolvability

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

The field of pervasive adaptation is an important area of modern research within the domain of adaptive systems. It is characterized by a dynamic collaboration between heterogeneous autonomous systems, complex real environments as well as multiple challenges of interactions between people, intelligent objects and computers. Robotics, which represents an essential part of modern technological society, is deeply involved into challenges of adaptive systems and shares research problems and solutions of pervasive systems (Kernbach, 2010).

Among typical problems, from such areas as adaptive control (Anderson et al., 1986), adaptive cognition (Keijzer, 2003), AI-based approachers (Bull et al., 2007), or softwareintensive systems (Wirsing et al., 2008), robotics involves into the domain of pervasive adaption new technological and biological inspirations (Floreano & Mattiussi, 2008). Considered conceptually, these inspirations and challenges are located around three following clusters:
— self-controlling, self-adaptive, self-learning and selfdeveloping of autonomous interacting system; control of such processes in a long-term perspective;

— increasing of complexity and appearance/emergence of undesired and less-controllable collective phenomena;
— development of bio-hybrid technologies and new principle of adaptation for systems with a high degree of developmental plasticity.

Control of long-term self-adaptive and selfdevelopmental processes. The issues of a long-term controllability of autonomous artificial systems is extremely important. Artificial adaptive systems with a high degree of plasticity (Levi & Kernbach, 2010) demonstrate a developmental drift. There are many reasons for this, like long-term developmental independency and autonomous behavior, emergence of artificial sociality, mechanisms of evolutionary self-organization. Such systems are very flexible and adaptive, but they also massively increase own degrees of freedom. New challenges in this area are related to a long-term controllability and predictability of "self-", principles of making plastic purposeful systems, predictability of a structural development and goal-oriented self-developing self-organization. These challenges have a great impact on a human community in general (the "terminator" scenario) as well as in different areas of embodied evolution, like synthetic biology or evolvable/reconfigerable systems and networks.

Increasing of complexity and emergence of social phenomena. Interaction networks in socio-technological systems can easily get so complex that the cause-and-effect chains are hidden and new collective phenomena are emerged. These phenomena are primarily related to appearance artificial sociality of technological artefact (Winfield & Griffiths, 2010), engineering the emergence (Stepney et al., 2006) and emergence of diverse "self-*" features. These phenomena are often very tough objects for control or even are undesired in many situations. Studying such phenomena, and in parallel, developing the tools needed to understand them is an important goal.

New principle of adaptation for systems with a high degree of developmental plasticity. Bio-hybrid systems, e.g. (Novellino et al., 2007) and mixed societies, e.g. (Caprari et al., 2005), are systems with a high degree of developmental plasticity. Neuro-robotic interfaces, in vitro neuronal networks (Miranda et al., 2009), syntectic biology (Alterovitz et al., 2009) are a few examples from these areas, which can completely change their own regulative, homeostatic and functional mechanisms. It is expected that bio-hybrid systems will have a large impact on ICT2 domain in the next 10-15 years. High developmental plasticity changes the postulated concept of "plant" and "controller", which originated from the control theory (Sastry & Bodson, 1989) and then widely used in other areas of adaptive systems. For these systems, new principles of a continuous and pervasive adaptation should developed.

References:

Alterovitz, Gil, Muso, Taro, & Ramoni, Marco F. 2009. The challenges of informatics in synthetic biology: from biomolecular networks to artificial organisms. Briefings in bioinformatics, November, bbp054+.

Anderson, B.D.O., Bitmead, R.R., Johnson, C.R.J., Kokotovic, P.V., Kosut, R.L., Mareels, I.M.Y., Praly, L., & Riedle, B.D. 1986. Stability of Adaptive Systems: Passivity and Averaging Analysis. Berlin: MIT Press.

Bull, L., Studley, M., Bagnall, A., & Whitley, I. 2007. Learning Classifier System Ensembles With Rule-Sharing. IEEE Transactions on Evolutionary Computation, 11(4), 496–502.

Caprari, G., Colot, A., Siegwart, R., Halloy, J., & Deneubourg, J.-L. 2005. Building Mixed Societies of Animals and Robots. IEEE Robotics & Automation Magazine, 12(2), 58– 65.

Floreano, D., & Mattiussi, C. 2008. Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies. MIT Press, Cambridge, Massachusetts.

Keijzer, F. 2003. Making Decisions does not Suffice for Minimal Cognition. Adaptive Behavior, 11(4), 266–269.

Kernbach, S. 2010. Towards Application of Collective Robotics in Industrial Environment. In: Rigatos, G.G. (ed), Industrial Systems: Modelling, Automation and Adaptive Behaviour. IGI Global.

Levi, P., & Kernbach, S. (eds). 2010. Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution. Springer-Verlag.

Miranda, Eduardo R., Bull, Larry, Gueguen, Franc¸ois, & Uroukov, Ivan S. 2009. Computer music meets unconventional computing: Towards sound synthesis with in vitro neuronal networks. Comput. Music J., 33(1), 9–18.

Novellino, A., D'Angelo, P., Cozzi, L., Chiappalone, M., Sanguineti, V., & Martinoia, S. 2007. Connecting neurons to a mobile robot: an in vitro bidirectional neural interface. Intell. Neuroscience, 2007, 2–2.

Sastry, S., & Bodson, M. 1989. Adaptive Control. Prentice Hall.

Stepney, Susan, Polack, Fiona, & Turner, Heather. 2006. Engineering Emergence. Pages 89–97 of: ICECCS 2006: 11th IEEE International Conference on Engineering of Complex Computer Systems, Stanford, CA, USA, August 2006. IEEE.

Winfield, Alan, & Griffiths, Frances. 2010. Towards the Emergence of Artificial Culture in Collective Robotic Systems. In: Levi, P., & Kernbach, S. (eds), Symbiotic Multi-Robot Organisms Reliability, Adaptability, Evolution. Berlin: Springer.

Wirsing, Martin, Banatre, Jean-Pierre, Hölzl, Matthias M., & Rauschmayer, Axel (eds). 2008. Software-Intensive Systems and New Computing Paradigms - Challenges and Visions. Lecture Notes in Computer Science, vol. 5380. Springer.

Pervasive Tracking in the Operating Room

Tamás Haidegger

Emergence and Evolvability

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

impacting (10) Importance

Computer-integrated surgical (CIS) systems strive to make medical interventions less invasive, less risky, less costly, but more efficient and capable of achieving better patient outcome. This requires providing clinicians with more useful information, prior to and during procedures, and more precise instrumentation. Creating modular, integrated systems comprised of imaging, computing and sensing, coupled with robotic devices and human-machine interfaces will radically extend the deployment of surgical robots. CIS systems are enabling clinicians to perform procedures with skills that before only a few had, and will open space for treatments currently out of reach.

Surgical trackers can localize tools and visualize patient's preoperative records in real-time. Navigation, as an enabling technology, has facilitated surgical simulation and training of medical students. In the future, the rise of the concept of "pervasive tracking" is expected, where all major parts of the medical working environment are tracked and modeled. It will be done real-time in a virtual environment according to their physical state and position.

One of the emerging modalities, electromagnetic tracking (EMT) already makes it possible to assist complex interventions, such as bronchoscopy, heart ablation and beyond. In the future, tracking technology should get embedded in the entire operating theater through the combination of EMT sensors and

RF ID tags. A fully sensorized operating room (OR) would allow for the better understanding of the procedures, the motion of the tools and the devices, track every movement of the patient, or the posture and the gestures of the surgeon- all without constraining the intervention. This will lead to an enormous improvement in interventional simulation and workflow analysis, and will result in the dramatic improvement of assistive surgical technologies.

Once we have a more complete understanding of the OR environment, the next transformational step in CIS is to advance the focus from enabling technologies to ubiquitous applications; extending the range of possible operations and the location of treatment. Surgical robots, teleoperated from hospitals, offices, or anywhere else by a skilled surgeon provide great advantages, and further resources could be freed if treatments were applicable beyond the OR. Providing an alternative to hospital- treatment would also mean bringing down costs, and opening new opportunities to health care related technical and practical innovation. To make telesurgery practical and accessible, fundamental changes are required in technology, medicine and health care management.

IV — Societies of Artefacts

Pervasive Systems Engineering

Achilles Kameas

Societies of Artefacts

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

challenging (15) Importance

It is now evident that pervasive computing is not just a new ICT research direction, but a novel technological paradigm affecting all aspects of human society. After developing small to medium scale systems, which operate efficiently in closed environments, I think that we are now in a position to consider how we can engineer the new generation of large scale autonomous pervasive systems. Contemporary pervasive systems are not modeled as systems, but as compositions of artifacts or orchestrations of services, following top-down engineering methodologies that place emphasis on functionality, resulting in systems only capable of responding to requests by using their constrained resources.

The research agenda item I propose concerns the engineering of autonomous intelligent systems that co-exist with people in real and synthetic environments (also referred to as "digital formations" or "digital spaces"). I anticipate the development of methodologies, processes and tools to engineer self-evolving pervasive systems composed of a large number of interacting autonomous individual artifacts, including robots, which can form societies, either long term or in opportunistic interest-based ad hoc manner, and can collaborate with humans to deal with situations that arise in shared environments. These individuals are engineered to intrinsically exhibit systemic properties, that is, to develop identity and mechanisms to preserve it, develop a sense of self in contrast to the "others", perceive the environment and pursue mutual adaptation, optimize their performance and deal with their shortcomings. Moreover, they are engineered to form societies, interact and compete with other ecologies, collaborate with humans and develop their own methods of conception and social norms. In the long term, I would expect such digital systems to exhibit properties of artificial life, be equipped with reproduction and optimization mechanisms, although they shall not be similar to those of biological systems, and ultimately draw their own evolutionary path.

This vision goes beyond the disappearing computer, cognitive robotics, complex systems and Internet of things and combines research results from ubiquitous computing, autonomous systems, cognition, human computer interaction, artificial intelligence, robotics, systems biology, genetics, sociology, complexity and other disciplines, into a novel engineering discipline, which aims to engineer systems that can develop their own internal structure, self-preserving mechanisms and structural coupling to the environment they operate in, including people. The ultimate goal would be to engineer systems that can develop their own representation of their world (their own "noema"), able to define a spiral course towards higher internal complexity and achieve evolution.

Energy efficiency in cloud computing

Albert Zomaya

Societies of Artefacts

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

An important technology that will enhance opportunities for adaptation for pervasive technologies is that of Cloud Computing. Cloud Computing is a paradigm in which information is permanently stored in servers on the Internet and cached temporarily on clients that include desktops, entertainment centers, table computers, notebooks, wall computers, handhelds, sensors, monitors, etc. This paradigm is gaining momentum rapidly with many companies offering cloud computing services (e.g. Amazon, eBay, Google, Microsoft) [1]. These systems are expected to dominate the way we do business in the next few years.

However, an important issue that will have great impact on how pervasive clouds will become is that of energy consumption. This issue in cloud computing systems raises various monetary, environmental and system performance concerns. A recent study on power consumption by servers (the type used in cloud computing systems) shows that electricity use for servers worldwide - including their associated cooling and auxiliary equipment- in 2005 cost 7.2 billion US dollars. The study also indicates that this electricity consumption in that year had doubled compared with consumption in 2000 [2]. Clearly, there are environmental issues with the generation of electricity. For example, the number of transistors integrated into today's Intel Itanium 2 processor reaches to nearly 1 billion. If this rate continues, the heat (per square centimetre) produced by future Intel processors would exceed that of the surface of the sun; this implies the possibility of worsening system reliability, eventually resulting in poor system performance.

As a result, there are efforts worldwide to design more efficient resource allocation protocols for these cloud systems. These protocols normally govern how an application or a set of applications are mapped onto the existing resources in a cloud environment. The design of such protocols for cloud systems is further complicated by the fact that the different sites in a cloud system are likely to have different capabilities. As such, it is not enough to simply minimize the total energy usage in the cloud; instead one needs to simultaneously minimize energy usage between all the different providers in the clouds. Apart from the multitude of ownerships of the different sites, a cloud differs from traditional high performance computing systems in the heterogeneity of the computing nodes as well as the communication links that connect the different nodes together.

The development of energy-conscious and power-aware resource allocation protocols for cloud computing systems will open up more opportunities for the deployment of more pervasive technologies. These protocols should lead to environmentally friendly "green" cloud systems whilst at the same time deliver the required performance levels (i.e. Quality of Service).

References:
[1] Carl Hewitt, "ORGs for Scalable, Robust, Privacy-Friendly Client Cloud Computing" IEEE Internet Computing, vol. 12, no. 5, pp. 96-99, Sep/Oct, 2008.
[2] J. G. Koomey, Estimating total power consumption by servers in the U.S. and the world, 2007.

Fluidity of Scale and Boundary

Ben Paechter

Societies of Artefacts

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Pervasive Adaptation (PerAda) refers to massive-scale pervasive information and communication systems which are capable of autonomously adapting to highly dynamic and open technological and user contexts. PerAda systems are thus a special case of collective adaptive systems which have particular constraints e.g. they are networked and highly distributed; they involve interaction with humans; they are large scale; the boundaries of systems are fluid; their context is dynamic; and they operate using uncertain information. We further note that they can be thought of as fractal in nature, in that a system consists of many subsystems, each of which can be thought of as system in itself such that each part is a reduced size copy of the whole with respect to the general challenges of adapting to the environment and in that the boundaries between the subsystems are fluid: in the horizontal dimension, at any time a system (component, devices, group, network etc.) may be joined by one or more other systems to form a new system incorporating the goals and context of all. Similarly, a system may split into two or more systems each requiring goals and context based on the parent system. In the vertical dimension a system might be considered to be at different "levels" depending on the viewpoint. Something that is a device from one point of view might be a component from another point of view. What we might view as a complete system (e.g. all the components, devices, groups and networks that exist within our car) might be seen as single component from another point of view (e.g. when considering the traffic flow through a city). Adaptation takes place across multiple scales; the adaptation process repeats itself from the small scale at the innermost level of the system (e.g. a component in a car) to the large scale outermost level (e.g. the traffic in a city).

Grassroots R&D, Prototype Cultures and DIY Innovation: Global Flows of Data, Kits and Protocols

Denisa Kera

Societies of Artefacts

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
impacting (10) Importance

Future belongs to innovation that simultaneously and directly connects politics with design, community building with prototype testing, and offers an experimental setting for following the impact of emergent technologies on society. This trend is embodied by the global rise of alternative R&D places existing outside of the government funded universities or even corporate R&D labs. In them innovation is becoming an active expression of citizenship as much as it is a human pursuit to understand nature and create resilient and efficient tools. Hackerspaces, FabLabs, Makerspaces, DIYbio labs, Citizen science projects based on Participatory Monitoring and Crowdsourcing of Data represent these alternative approach to R&D that combines decentralized approaches to management and policy with P2P, open science and open innovation approaches based on open source hardware and software infrastructure. These social as well as technological experiments and prototypes define new niche markets and enable global and sustainable exchanges that have a potential to empower various communities. Global and alternative innovation networks are developing around Do-It-Yourself (DIY) and Do-It-With-Others (DIWO) subcultures, such as Direct to consumer (DTC) genomics, DIYbio labs, DIYgenomics, Clinical trials 2.0, Hackerspace hackathons, Maker fairs and FabLabs competitions. Communities of people monitoring, sharing and making sense of various scientific data and technological practices are exploring these new global networks around low-tech DIY and open science protocols. Maker and hacker communities around the world prototype future gadgets and tools with open hardware platforms and feed the needs of various grassroots open labs for affordable equipment that offer opportunities for entrepreneurship. These low-tech and open source strategies are paradoxically inspired by both EU alternative squat cultures and the American spirit of entrepreneurship. The global and alternative R&D places are made possible by informal networks around the globe that enable very different flows of knowledge and expertise from the official industry and academia. They are becoming testbeds for new models of public participation in Science and Technology but also new models for policy making in which political deliberation merges with design iteration and embraces citizen science paradigms of research. We are entering an age of global flows of open data, DIY kits and citizen science protocols.

Theoretical Foundations of Socio-Cognitive Systems

Jeremy Pitt

Societies of Artefacts

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

It is almost tautological to state that "societies of artefacts" should be based on social relationships, but this a challenge yet to be fully addressed. It requires an inter-disciplinary investigation by social and computer scientists to explain how social and socio-cognitive relationships between people can be represented in a computational theory and used as a basis for engineering societies of artefacts. In particular, this research programme should investigate the principles of social intelligence and social networks that underpin the specification and regulation of open, decentralised, and adaptive systems, and provide corresponding mechanisms for run-time self-regulation and self-management. This will involve study of the innovative intersection of norm-governed systems, voting algorithmics, game theory, opinion formation, belief revision, judgement aggregation, and social computational choice [PKSA06], as well as a formal characterisation of socio-cognitive principles of trust, forgiveness, and affect. It will also need to leverage results from the mathematical study of social networks: a particular area of study here is the interaction of aggregation rules, social networks, and intended outcomes, for example in reaching consensus, avoiding dominance, maximising utility, minimising inequality, and so on [Dra06, DG09]. A successful outcome would serve to establish the logical and computational foundations of a common theoretical framework for the executable specification with respect to runtime "social services" provision, as a basis for performing secure and reliable service discovery, composition, optimization, adaptation and deletion.

References:

[DG09] M. Draief and A. Ganesh. Spread of epidemics and rumours with mobile agents. In European Consortium for Industrial Mathematics Symposium (ECMI 2008), to appear, 2009.

[Dra06] M. Draief. Epidemic processes on complex networks. Physica A: Statistical Mechanics and its Applications, 363(1):120-131, 2006.

[PKSA06] J. Pitt, L. Kamara, M. Sergot, and A. Artikis. Voting in multi-agent systems. Computer Journal, 49(2):156-170, 2006.

How to program populations of tiny artefacts

Paul Spirakis

Societies of Artefacts

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Pervasive environments of tomorrow will consist of populations of tiny artefacts (passively or actively moving perhaps) and interacting with each other. The state of the art in designing and implementing protocols for such populations is still at its infancy. Programming happens at a low level, and the programmer has to deal explicitely with the peculiarities of the system calls supported by the tiny devices. Debugging is not easy and actual experiments of large scale are needed.

Thus, a big challenge is how to program such populations in the large. For example, one would like to state high level "suggestions" like "reduce your energy spending", "merge two populations", " please elect a leader", "spread the information by an epidemic process", "increase the security level" etc and, ideally, the underlying population should be able to implement these in a scalable (independent of current population size) and flexible way (e.g. choose among a variety of routing methods). The issue becomes more challenging when the populations are "selfish" ie when devices have individual objectives! Then, the high level language may consist of "motivations" and "suggestions", reminding the use of advertizements in markets or the use of financial motives and rewards! Such a challenge needs an abstract model of pervasive populations and high level language constructs with well defined semantics. We note that such a high level programming ability has to be accompanied by significant progress in the direction of an "intelligent middleware" of populations of tiny artefacts.

This challenge requires progress in distributed computing, game theory and economics motivated research, new logics and semantics, and even new notions of controlling the behaviour of such "social systems" while hiding cumbersome low level details (heterogeneity, multiobjectives of devices, detailed local communication, nature of interactions, even device identities and population size). Note that each such high level command or suggestion activates a whole dynamic process in the population, and its "final" result is not instant. In fact a wrong use of a sequence of such "commands" may raise even chaotic behaviour in the underlying populations. Stability notions, and transient behaviour may require a blend of mathematics not much used by the Computer Science community (eg nonlinear dynamics or stochastic processes). The gain will be enormous with respect to design of many diverse applications in the pervasive environment. Notions of AI may be useful here (in fact, of distributed Artificial Intelligence), and useful paradigms may be drawn from biological groups of animals.

Content and Communication without the Internet

Roger Whitaker and Stuart Allen

Societies of Artefacts

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

The Internet has now brought electronic content and communication to society in an accessible form. It allows content to be accessed and shared through web-based technologies using a diverse range of different devices, which are increasingly mobile. This is very much a first phase in providing electronic communication and content to humans, with the defining characteristic being the use of backbone networks that facilitate content delivery.

However this is not the only way of providing content. Hardware is becoming increasingly more portable, cheap, powerful and short-range unlicensed wireless communication standards are also emerging, allowing diverse devices to access and share data with their peers through transient interactions when they come into range with each other. This is being enabled through the development of opportunistic networking architectures. Providing content and communication using these means is not a replacement for the Internet. Instead, new applications and services are likely to be facilitated that permit mixing of digital content with physical spaces and enhancement of human social communication and interaction.

This results in significant research challenges, based around important emergent trends. The participatory generation of content is now a common phenomenon that has led to the concept of the prosumer, where network users are both producers and consumers. Such content can potentially be directly acquired through mobile peer-to-peer networking between mobile devices opening up opportunities for social mechanisms to propagate information without the Internet through human social networks. Examples of future applications may include opportunities for further human self-expression, such as the use of micro-blogs, which can be spread through person-to-person interactions to propagate knowledge.

Further applications may also involve previously inanimate objects, appliances, products and environmental "furniture" that are able to share information about resources and services. This embodies the vision of pervasive and ubiquitous computing and also emphasises the need for development of geo-informatics to support spatial content provision. Ultimately, this vision will lead to a physical world that is saturated with content, far beyond the current Internet, and will require new techniques to recognise and acquire content that is of relevance to situation or context. Unlike the conventional Internet, pulling content is more challenging because transient mobile peer-to-peer interactions are naturally push disposed. This also means that centralised search engines are unable to provide indexing of content and its relevance as for the Internet. Consequently new autonomic paradigms are required that can navigate a large content space and match it to users and situation. This is a long-term pursuit, the foundations of which can provide far-reaching functionality.

V — Dependable Pervasive Systems

The challenging fact about pervasive adaptation is that our current general purpose software and systems engineering methods do not fit to the task of designing and implementing pervasive systems. Quality in systems design and in software engineering is achieved today by methods like formal engineering, strict design processes and intensive testing. However, these methods require a specification to verify, design or test. The inherent property of adapting pervasive systems is that such a rigid specification cannot exist a-priori as the systems need to adapt to unforeseen circumstances, both in the context of their use and in the systems they may need to cooperate with in the future.

A number of novel general-purpose programming models have been proposed for pervasive systems, ranging from classic distributed systems and service-based models with additional properties supporting pervasive applications to programming models based on an ontology and inference rules dealing with specific properties of pervasive systems such as context awareness. Many special-purpose software systems have been proposed to deal with specific aspects such as distributed sensor data processing or user interface concepts for pervasive systems. Unfortunately up until now, no general programming model has emerged that could support a reliable large-scale and long-term deployment of pervasive systems consisting of components of various origin and original purpose.

Two essential requirements for such a future programming model are support for adaptation and verification. Adaptation provides the mechanism for a system to adapt to new circumstances, e.g., to new environments, to changing environmental conditions, to new users and tasks. Verification provides the mechanism to trust a system. Through verification, the system can detect circumstances in which it cannot fulfill its purpose, i.e., it cannot provide results or services according to functional or extra-functional requirements. Verification can also provide proof that, given the initial assumptions are correct, the system will maintain certain properties.

To enable both adaptation and verification, pervasive systems must provide other systems with a formal description of their current requirements and

Sustainable Software Engineering For Pervasive Systems

Holger Kenn

Dependable Pervasive Systems

Industrial Contribution

medium (0) Agreement

long term (>3 years) Duration

challenging (15) Importance

properties. Based on this information, other systems can decide if and how they can cooperate with each other. This information forms a living specification of the global system that evolves at runtime and can be used to infer emergent system properties.

Norms for Dependable Systems

Jeremy Pitt

Dependable Pervasive Systems

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

challenging (15) Importance

Any networked operation with a dynamic topology that features distributed functionality, decentralised command and control, local decision-making, delegated responsibility and autonomous action can be regarded as an instance of a normative system. A normative system refers to any set of interacting agents whose behaviour can be regulated by norms, where a norm is a rule which prescribes organizational concepts like power, permission, right, duty, sanction, and so on [Ser01].

Teams of robotic vehicles, in military, search-and-rescue, disaster-recovery or surveillance missions, operating under no, partial or even full human remote control, undertaking missions of high complexity and uncertainty, provide just such a (inter-)networked operation and exhibits precisely these features. Moreover, these teams have to be self-regulating, in terms of a dynamic re-allocation of roles, tasks, priorities etc., which can be specified as part of the normative system itself. A major challenge is to define dependablity, in terms of being able to meet specific organizational objectives and levels of utility, at the same time being able to withstand component-loss, network outage or overload, and/or hostile behaviour. This would entail, one the one hand a detailed investigation of the logical and computational foundations of normative systems for self-regulating dynamic networks of teams of robotic vehicles that conduct missions in potentially hazardous environments; and on the other new models of event recognition from massively parallel data streams and distributed sensor networks.

The proposed research can provide organizational benefit to e.g. emergency or military services by developing a new paradigm for dependable operations which:
(a) is consistent with the trend towards fewer, but smarter, mission components capable of assimilating and acting upon local intelligence 'on the fly';
(b) expands the range of tactical mission options through increased operational flexibility and resilience to attrition;
(c) supports coordination across service boundaries, resulting in seamlessly integrated joint service operations; and

(d) through the formalisation of normative concepts, which are essentially human constructs, supports modeling or maintenence of the ethical judgement and moral imperative of a human operator.

References:
[Ser01] M. Sergot. A Computational Theory of Normative Positions. ACM Transactions on Computational Logic, 2(4):522-581, 2001.

The Internet-of-Humans

Max Mühlhäuser

Dependable Pervasive Systems

Academic Contribution
medium (2) Agreement
long term (>3 years) Duration
challenging (16) Importance

As computers become truly ubiquitous, the Internet as the global integration backbone proliferates: 1. The Internet-of-Things (IoT) hosts networked everyday objects and the "six senses" of the global ubiquitous computing infrastructure, 2. The Internet-of-Services (IoS) comprises the ever-growing world of software applications available, increasingly modular, interoperable, and self-organizing - in the future, this facet of the ubiquitous Internet will inhale more and more of the world's economy as a key globalization driver; 3. The Internet-of-Crowds represents the power of Web2.0 and social networks, currently still rather disjoint from both IoT and IoS, but with a strong tendency to blend with them; 4. The Internet-of-Clouds coins the transformation of the 'underlying' Internet - in the past basically a universal, global 'bit pipe' - into the general-purpose workhorse for the aforementioned three Internet facets; in this role, the Internet-of-Clouds will provide pervasive storage, computing, and communication resources accessible to anyone anytime in any amount; an open question remains the extent to which cloud resource providers (including search engine providers) will or should be - or can be prevented from - possessing the footprints and fingerprints of all of the world's citizens and organizations.

The above-mentioned four facets of the ubiquitous Internet are well known and widely discussed. This contribution, however, aims at highlighting a fifth, probably often `forgotten' variant of the Future Internet: the Internet-of-Humans. In order for the IoT and IoS to happen on a global scale, Humans must be empowered with a service-independent ubiquitous access to the Future Internet. This access and use concept has to hold and enforce the user's wishes and needs in a fully trusted way, independent of any 'provider' but compatible with all the services and 'smart things' encountered on the globe.

The use of an ultra portable, personal device (for which the cell phone might be a crude starting point) is advocated as the individual's "key" to all other

Internet facets. This 'key device' should realize, among others, the two most important aspects of the Internet-of-Humans:

a) an utmost level of usability, hosting and mediating between the abstract (!) user interfaces of all services accessed, turning them into concrete physical user interfaces and leveraging multiple modalities and interaction devices as dynamically accessible in the user's vincinity;

b) maximum privacy, confidentiality, and trust as a key enabler to acceptance of pervasive computing power and 'big brother potentials'. For instance, the aforementioned key device should host the core of a pseudonymity/identity/authentication/encryption mechanism, strictly enforcing the user's privacy and security desires; it should also hold and secure indispensible keys required in order to interpret the user's "footprint and fingerprint" information in the above mentioned sense (cf. Internet-of-Clouds).

Building and Assessing Dependable, Usable, Secure and Evolvable Safe Pervasive Systems

Philippe Palanque

Dependable Pervasive Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

Pervasiveness is a property that most computer systems are attempting to offer and it is very likely that such trend will carry on in the future. One of the issues raised by pervasive computing systems is that they gather many attributes that are unlikely to be otherwise present at the same time in computing systems. Namely they are typically interactive, collaborative, embedded and adaptive and are likely to be used in safety critical context (healthcare, aeronautics ...). This require to integrate research contribution from different fields including secure and dependable computing, human-computer interaction, groupware and social computing and safety engineering. The research questions will thus be at the intersection of previously more independent research field. Their nature will additionally raise new research questions due to their adaptive nature. Indeed, such requirement will impose on the designers of such systems to take into account the various possible contexts of executions and to define and built adequate designs for each of them. In addition to this design/construction problem the issue of assessment of the quality of these adaptive systems will be of primary importance and will for sure be a challenge to be addressed if they are to be deployed in real, large scale and safety critical contexts.

Large-Scale Impact of Pervasive Computing

Shwetak Patel

Dependable Pervasive Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Discovering ways for us to have high value impact today remains a challenge in the Pervasive Computing community. We have struggled to enable large-scale explorations of socially meaningful applications. These applications include home health, elder care, and energy and resource monitoring. Often in research, solutions are limited to the laboratory and small deployments because of the capital and labor required for their installation and maintenance. To reduce this installation burden, we need to work on new sensing techniques and approaches that can start to take the leap out of the lab and into real homes, focus on non-expert installation and provisioning of these devices, and engage with other domains that can help advocate the value of such large-scale deployments. We have already seen this trend in a variety of research efforts that are beginning to leverage existing infrastructures and commodity devices to quickly bring their technology to the public. Clearly this concept is not new, but it is worth advocating that the research community needs to embrace that there is still real science in making technological solutions practical and available. It is not until we start to make this leap forward, will we start to have the same level of impact that other disciplines have seen.

Modeling of Pervasive Systems and Real-World Processes

Stephan Haller

Dependable Pervasive Systems

Industrial Contribution
medium (0) Agreement
mid term (>1 year) Duration
impacting (10) Importance

With pervasive computing, it not only becomes possible to collect data from the real-world, but also to execute some form of business logic in a distributed fashion, close to the "point-of-action" in the physical world. At its extreme, this leads to smart and autonomous physical objects that collaborate to a certain degree in order to achieve specific goals. It can be argued that such distributed execution and autonomy is needed to achieve the required degree of scalability, the desired response times when interacting with such objects, as well as the resilience against potential network failures. However, how can it be guarantueed that a system of autonomous objects is doing what it is supposed to be doing? How to deal with the inherently unreliable aspects, like devices suddenly failing or decalibrated sensors? And related to that, how do I manage such a system? An answer to these questions is crucial for any organisation - in particular commercial enterprises - that thinks about deploying and utilizing pervasive systems.

A suitable approach to overcome this major roadblock would be if such systems could be modeled - the overall system, as well as individual objects and resources, and any constraints. This would allow formal verification as well as simulation, and with the development of appropriate execution platforms, also the manageability would be improved considerably, e.g., by deploying functionality on the fly if the model is changed. Concretely, the following issues must addressed:

— Platforms allowing the easy integration of pervasive real-world services into overall process modeling, in particular business process modeling. Currently integration of such services require too much specialized know-how, but the inclusion of real-world aspects into business process modeling tools would significantly increase the adoption in industry and allow the creation, deployment and execution of sensor-based processes.

— Models for deployment and operation of distributed pervasive systems must be found that are sustainable in the sense of cost, life-cycle management, energy-efficiency etc.

— Measures for the Quality of Information (QoI) of data from sensors and sensors networks are essential if any decisions - in particular, in M2M scenarios - should be based on that data. It must be possible to assess the QoI not just of individual sensor data, but also of the fused data from several sensors and any (complex) events derived therefrom.

VI — Pervasive Trust

Nowadays, the pervasive deployment of tiny devices with minimum storage and limited or no computational capabilities appears a realistic perspective. We refer to a new class of passive devices (i.e. not powered by batteries) that has emerged in the last decade, the most prominent examples being RFID and NFC tags and new ones that will appear in the next years.
These devices are carried by humans and are able to collect and distribute data collected both for social and private use. Plenty of applications are foreseen impacting the society (e.g. monitoring pollution or traffic congestion monitoring information obtained by citizens? devices) and individual (e.g. personalized recommendation in shops, information concerning the individual social network, health monitoring).

A crucial issue in the effective deployment of such applications is to convince users to allow the delivery of the information that is collected; in fact, the exploitation of such devices in order to perform suitable operations should guarantee privacy of the user and trust that devices are not compromised.

Here we focus on the privacy issue. As an example an individual might be interested to deliver to the city traffic agency the information concerning the traffic congestion she has encountered while driving, but at the same time she does not want to reveal her destination neither to other users nor to the city agency itself.
Collecting data while guaranteeing user privacy requires an interdisciplinary approach and new ideas from several areas. In fact, the solution should be compliant with the limited resources of tiny artefacts (both computational power and memory and energy) and should be transparent to the user (i.e. not requiring user intervention).
The effective solution of such issue will allow to fully exploit the data deluge capabilities of pervasive devices thus allowing the deployment of systems and applications with added social value.

Privacy preserving cooperation in pervasive scenarios

Alberto Marchetti-Spaccamela

Pervasive Trust

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
pragmatic (5) Importance

Pervasive Trust

Eve Mitleton-Kelly

Pervasive Trust

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

The European project SOCIONICAL is developing and testing an app for smart phones, to both elicit data during an event where large crowds are involved, as well as to provide information to the user. In experiments, the data gathered relates to density of crowds, direction and movement; while most of the information offered to the user will be about the sporting or social event. In the event of an emergency, the relevant authorities would be able to advise users on the best (least congested) exits to use or give other directions on evacuation.

The critical question however is, would users trust that information? Why? Would they act on it? If they know the provenance or source of the information and trust that source, then the user is likely to trust the information and act on it. But this is merely an assumption.

In an emergency, trust would also be involved in accepting help or advice from a complete stranger. If your life depended upon it, why would you trust someone you knew nothing about, except that they were offering help or advice?

We know that trust is not a simple relationship and that there are different types of trust. Also that risk is required for trust to influence choice and behaviour; and risk is generally related to vulnerability and/or uncertainty about an outcome. A working definition might be "Trust as a willingness to rely on another party and to take action in circumstances where such action makes one vulnerable to the other party."

In an excellent paper Doney, Cannon & Mullen (1998) identified five cognitive trust building processes, based on cultural norms and values that facilitate the application of each trust building process. They are: (1) calculative process: trust is established when one party calculates the costs and/or rewards of another party cheating or cooperating in a relationship; (2) prediction process: trust is based on prior experience based on the assumption that behaviour is predictable; (3) intentionality process: assumes a link between intention and trust and is based on the notion of benevolence; (4) capability process is based on technical competence as a precursor for trust; (5) transference process: whereby trust is transferred from a trusted "proof source" to another individual or group with which the trustor has little or no direct experience, e.g.certified public accountants may be trusted because they are certified by a trusted agency.

The five cognitive processes provide a framework of trust-building processes and suggest that a trustor's willingness to act is based on trusting expectations, which develop through a pattern of cognitive analysis. However these processes build over time and we still do not know the cognitive process of trust in an emergency, when decisions have to be made quickly with no prior knowledge or experience of the other person. Nor do we know why information provided through an AmI device could or should be trusted. These are critical questions which need further investigation.

VII — Human-Centric Adaptation

Within the framework of the EU-funded project ATRACO (Adaptive and TRusted Ambient eCOlogies) our research focuses on uncovering "what comes next" within the Area of Intelligent Environments. Starting from the definition of Ambient Ecologies (AE) that conceptualize spaces populated by interrelated and therefore connected devices and services, various environments and people we establish the term Symbiotic Ecologies (SE) to describe what happens when human and AE collaborate and benefit from each other. A generalization of SE is what we call Next Generation Ambient Intelligent Environments (NGAIE). NGAIE support the realization of Activity Spheres (AS) that are modelled as dynamic and purposeful projections of SE.

ATRACO aims at the realisation of AS by focussing on the role of knowledge within such realisations. Knowledge bridges the gap between the levels of a layered architecture forming an AS. From top to bottom we define the following levels: HCI, Semantic, Behaviour, Structure, and Communication. These levels also describe the facets of concreteness and abstraction towards user with respect to his perception of the system. On top the system is concrete from a human point of view and abstract from a system point of view. As the user shifts focus towards the bottom, the system components become concrete, but it operation becomes vague.

Heterogeneity is an intrinsic NGAIE feature that causes problems in practical AS operation. For each level we recognize a specific type of heterogeneity. For example HCI realised by a spoken dialogue system is confronted with the heterogeneity of human speech. How to handle this heterogeneity is still an unresolved research issue. Ontology alignment provides a solution for representing heterogeneous knowledge on the Semantic level. Advanced Fuzzy technologies can be used to adapt the behaviour of the AS to the heterogeneous behaviour of the user. On the structural level ontologies can be used to provide knowledge about the changing population of entities within the AS and finally on the communication level adaptive middleware and system hardware can be used to maintain the technical interrelation of AS entities.

The ATRACO approach combines many research topics within the area of

Symbiotic Ecologies in Next Generation Ambient Intelligent Environments

Achilles Kameas and Tobias Heinroth

Human-Centric Adaptation

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

challenging (15) Importance

Intelligent Environments and integrates their results in small scale systems that support AS. The next step is the realization of large scale symbiotic NGAIE. Scaling up is a big challenge, as it affects the complexity and performance of all system layers and the interactions among components. Indicative challenges include dealing with large scale semantic heterogeneity in near real time, dynamic formation and dissolution of AS, multiple loci of interaction and control, concurrent activity recognition, resolution of conflicts, adaptive trust and privacy policies, different networking and communication protocols, etc. In addition, symbiosis of AS with humans raises issues such as autonomy, perception, trustworthiness and dependability and even developing higher order capabilities, such as learning, assuming responsibility, taking initiatives and dealing with uncertainty.

Approaches for Context Creation and Handling

Adam Wolisz

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

There is a clear trend to develop and offer more and user-personalized services. These services are dependent on elaboration and storage of a broadly understood context: Information about the user, his habits, his location, his activity, history of his usage of the service e.t.c.

The problem is, that nobody is aware which application, which provider is storing what, nobody is in real control of the information stored about him, nobody can anticipate how and when this information might be used. On the other hand, the advantages of "context based applications" push users to: register, answer questions or click "I agree" without reading numerous pages of cryptic information.

There is a fundamental need of developing new information processing paradigms. We have to develop the approach of "stateless context based services" - in which the service provider is NOT allowed to store ANY user context related information, but in contrary, each individual query is enhanced - if needed! - by the desired context information.

Where would such information come from? There is a need for developing concepts and toolset for "unified profile description and profile usage". Such toolset might be downloaded by a user and maintained by him, or - more frequently - the user might create a set of "profile accounts" with a selected, trustworthy, "profile service agency" (like we have one or more bank accounts today!).

Each query for a context based service could be bound to one such account (possibly having different level of information!). The results of the query might - if so permitted! - even update the profile. This would be - for example the case if somebody agrees to store in his profile the history of used services or history of evaluation of the satisfaction with the service... etc.

The crucial issue is that the content of the profile would be always stored in well known, trustworthy entity and controllable (including the delete option!). Each service requesting context data would have to explicitly define in it interfaces which data are requested, and carry a legal commitment not to store any context/personalized data provided in a query. The ("trusted") profile service agencies could offer advice on recommended level of profile complexity, and "depth of context information" to be opened to services - similarly as we expect banks to give us risk assessment for investment (and - yes, we know that this did not always work :))

Such approach offers a conceptual framework to assure the desired level of privacy - but also the flexibility needed for context based services.

We need a pretty elaborated set of concepts for such profile description - and interaction pattern. It has to be flexible enough to allow adding of new attributes, and new ways of history storage and processing. It has to be bound to legal acts.

But it opens a perspective to deal with a fundamental problem; the challenge is simply there and can not be ignored...

Universal Proactive Interface to Pervasive Services

Anatole Gershman

Human-Centric Adaptation

Industrial Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Pervasive computing increasingly penetrates our physical environment, both private and public, from smart appliances to buses, offices, and coffee shops. Each instance offers some kind of services: comfort controls, schedules, prices, music, etc. We fully expect a dramatic increase in such pervasive but location-specific and object-specific services. Today, most of these services are accessed through unique physical interfaces such as vending or airport check-in machines, or your toaster control panel. With the proliferation of such services, many of them will be offered wirelessly and accessed through the one truly personal computing device all of us are carrying - the mobile phone. This will make the phone a "universal remote control to the world" which raises the first question: what sort of common controls might emerge? Will we see the equivalents of the "file, edit, view" menus familiar to Windows users?

Today, when I pull my mobile phone out of my pocket, it always shows the same picture of my favorite sunset whether I am running through the airport trying to catch my plane or standing on a street corner at lunch time. I have to invoke specific applications to find out if my plane is late or where to have my sashimi. Yet, most likely, my smart phone knows where I am, what's on my schedule, and my preferences, which should be enough to proactively give me the information I need when I need it. What kind of intelligence would this require? We might start by analyzing some of the most common concern/context pairs: business traveler at an airport, office worker driving to work, teenager at a shopping mall, etc. Each pair would define the likely sources of the relevant information and the priorities for its presentation on a small screen.

A more sophisticated version of an automated personal assistant should continuously monitor the available pervasive services and decide what and when should be brought to the owner's attention. This must not require exhaustive specification of decision rules by the owner. The system must start with a great deal of general knowledge and then adapt itself to the owner by learning the patterns of his or her behavior. It must also be able to adapt itself to ever changing pervasive services and environments. In addition, it should be able to take explicit instructions from the owner. While creating such an assistant will require solving many hard Artificial Intelligence problems, even relatively simple versions will be very useful. Otherwise, pervasive services might be represented by a very long list of specialized apps on our phones or require service-specific devices.

Self-organising Knowledge-intensive Workspaces

Andrea Omicini

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
challenging (15) Importance

Knowledge is pervasive in the information age. Knowledge-intensive environment represents an essential dimension in modern pervasive systems, where rich information contexts and knowledge-oriented human-computer interaction are often fundamental elements in the design of pervasive applications.

However, the (too) many possible ways of human-computer interaction in pervasive systems, along with the huge, unordered amount of relevant information available could make the design of knowledge-intensive workspaces a nightmare. A computer-based working environment for researchers, journalists, writers, lawyers, physicians, politicians, administrators (all professionals dealing with critical amounts of information resources that need to be suitably reviewed, organised and used within complex, ever-changing contexts) requires on the one hand that all the relevant information sources are made available to the user in a complete yet usable format, on the other hand that the working environment autonomously evolves and adapts to the individual uses and work habits, as well as to the diverse goals and contexts defined by the everyday activity. Both requirements could be addressed by facing the following key issues: the suitable definition and representation of all the relevant information; the self-organisation (including creation, propagation and dissipation) of knowledge; and the self-organisation of the working environment.

While relevant information could be handled by adopting many well-known techniques in the area of knowledge representation and extraction, the main point here is the explicit representation, memorisation and exploitation of user actions in the workspace: user actions should leave traces on the working environment which (like in stigmergic coordination, or, more generally, as in the case of cognitive stigmergy) are to be used for the self-organisation of both the relevant knowledge and the workspace.

Furthermore, suitably-represented information chunks could autonomously combine to create new knowledge memes that could propagate towards the user (using for instance nature-inspired mechanisms like chemical or field-based ones), thus providing users with a self-organising knowledge-intensive environment. Self-organisation of knowledge would involve both the individual workspace (where it should be mainly used for the self-organisation of the working environment) and the information infrastructure, and would in principle affect all the available dimensions: individual, intra-organisation and inter-organisation, up to the global scale of the Web as a world-wide knowledge source.

Among the main challenges to be faced to this end, one could mention:

— understanding the best techniques for representation and exploitation of user actions in the working environment;
— devising the most effective mechanisms for knowledge self-organisation -- including knowledge creation, propagation and dissipation -- to be used both at the individual (in the workspace) and the global (in the information infrastructure) levels;
— defining the most effective architectures for the individual workspaces and the information infrastructure, which could enable and promote effective self-organisation within any sort of pervasive scenario.

Intelligibility in ubiquitous computing systems

Anind K. Dey

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

We have overcome some of the biggest challenges for doing ubiquitous computing research: lack of available tools to build applications, lack of available infrastructure to support applications and lack of end-user devices to foster widespread use of applications. Particularly with the recent boom in the use of smart phones and location-based services, we can see that many of these challenges have been addressed enough to allow for commercial success. However, now that we can build interesting and compelling applications, a remaining challenge is to make these applications usable and easier to adopt.

One particular usability aspect of interest is intelligibility, helping users to form an accurate mental model about how to use an application. This is important for allowing users to understand how the application works and to be able to predict what it will do in a future situation, and all of this will impact adoption and use. This is particularly true for new types of ubiquitous computing applications that involve adaptation or machine learning. In both of these settings, the level of complexity involved can be confusing to a user (e.g., recommender systems and Microsoft Clippy). When users have difficulty forming a mental model about how these work, they are less likely to adopt them and use them.

As a research agenda, intelligibility in ubiquitous computing systems is wide-reaching. At the lowest level, it requires the collection of low-level sensor information and data that may feed into a machine learning or adaptive system, from a wide variety of systems: basic context collection. Then, it requires inference or analysis of that data to achieve higher level understanding of the situation, possibly using machine learning. This output can then be fed into a machine learning or rules-based system to determine what action should be taken on behalf of a user. At each one of these stages, the computing undertaken must be somewhat transparent, to allow the generation of an explanation of what has occurred: what data has been collected and why, what inferences have been made and why, and what actions have been taken and why. This is a challenge when dealing with experts and machine learning systems, but even harder when dealing with new users and complex systems that involve sensors, inferencing, learning and actions, where some of the resulting actions may not be visible to the end user herself. Such intelligible systems must be developed, allowing users to ask why did the system take a particular action, why didn't it take a different specific action, how can I get the system to take a particular action, and what will the system do in a specified situation. It is only through the development of such intelligibility in ubiquitous computing applications, will we have widespread adoption and usability.

Lifelong human-machine co-adaptation for pervasive smart assistance

Daniel Roggen

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

The provision of pervasive smart assistance requires a conjunction of elements: sensing of the user's context, deriving the user's needs from his context, and activation of a support strategy to address these needs. Unobtrusiveness is highly valued. Context-aware pervasive smart assistants fulfil at first view the definition of an "adaptive" system, as the system's function is adjusted to the situation of the user.

The contexts of interest, the support strategies, and the mapping between them are usually defined at design-time using various heuristics, such as deriving them from training data, from observational studies, or from expert knowledge. Yet current systems fail to realize the full potential of an adaptive system, as these elements remain static throughout operation. To some extent, this limits the degree of unobtrusiveness that can be reached as it casts the need-to-support pattern in a predefined mold. User's requirements are likely to change over time due to aging, injuries, or changing interests.

We believe that the next frontier in pervasive smart assistance will be to devise systems capable of continuous - lifelong - co-adaptation to the user needs. Thus, the context of interest, the support strategies, and the mapping between them need to embrace this dynamics and become themselves adaptive. However, towards what should be the system adapted to, and by what should adaptation be driven? We believe there is no better way than to be guided by the satisfaction of the user when he is interacting with the system.

Concretely, next generation smart assistants will need to continuously evaluate the degree of satisfaction (or frustration) they elicit in users. They will then explore other assistance strategies, continuously assessing the outcomes of

their new actions, learning from their mistakes, and refining the better strategies.
Realizing this lifelong human-machine co-adaptation will require pervasive smart assistants to use all available resources to gauge the user's perception of the system.
Minimally obtrusive mobile interfaces will be needed to collect explicit user feedback.
Implicit user feedback will come from the machine recognition of emotional states such as frustration or satisfaction, likely using a wide range of sensing modalities, such as electro-encephalography, physiological and motion sensing.
New online learning methods will need to be devised to capitalize on all sorts of sporadically available and potentially noisy feedback information.
This will call for an effort of many communities including HCI, machine learning and obviously the pervasive computing community.

Creative Methodologies for Informing Design of PerAda Systems

Ingi Helgason

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

The pervasive computing vision is one where humans will interact with invisible technologies in a seamless, effortless and intuitive manner. This is a beguiling image of the future, one where our physical environment adapts to us, instantly, in a safe, immediate and pleasurable way. Our individual personal spaces and intimate belongings will know who we are, and what we need at any time of day, wherever we are. No longer will we be required to make our own decisions about what is good for us, as technology will know better. Our social relationships, healthcare, workplaces, homes and cities will all be affected.

It will be a substantial challenge to design novel systems that will have such a radically transformative impact on our lives. Complex and difficult questions will need to be addressed, and as researchers and as members of society, we do not even know with certainty what these questions will be. Increasing our understanding of the experiential, dynamic and subjective nature of human interaction with systems will be vital. Research into human factors such as the myriad contextual, social and situational influences, and the range of individual variations in people will also be paramount.

One approach that can be helpful in beginning to examine these issues is to collaborate with artists and creative designers in a critical exploration of fictional possibilities. Technologists already value the input of scientists from domains such as psychology, ethnography and sociology. These fields of study can provide insights that usefully inform the design and implementation of systems for maximum acceptance. Taking this multi-disciplinary approach further, by involving the creative industries, is also gaining acceptance as a way to inform scientific research directions. A recent example of this is the Human+ exhibition (http://www.sciencegallery.com/humanplus) taking place in Dublin at the time of writing. This exhibition brings together installation work by artists that considers the future of the human species in the light of scientific advances. By presenting extreme and intriguing possibilities within a publicly accessible context, debates and dialogues are initiated, and alternative views can be considered.

We all have hopes and fears, priorities, values and psychological needs that are fluid over time. Technologists must take these into account when developing systems that will be accepted by individuals, and beneficial to society. Issues of security, privacy and trust, in particular, are a priority in pervasive and adaptive systems. In previous stages of Human-Computer Interaction (HCI) research, conceptual models, user feedback and visibility have been at the forefront. When systems disappear, as in Weiser's vision, the user is set adrift, and a new design methodology is needed to enable the user to feel secure again.

Personal physiology monitoring system based on time-series analysis

Jan W. Kantelhardt

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Pervasive sensors placed at different locations on the human body or in the clothing can continuously monitor a large variety of physiologic parameters (time series), such as heartbeat, brain waves, blood pressure, oxygen saturation, muscular activity, respiration, body temperature, etc. If pervasive computers are embedded with such sensors they can jointly evaluate these various data streams in real time, extract relevant features, and classify the state of the human body not just regarding activities or motion patterns but also regarding specific medical conditions and the healthiness of the current lifestyle. This way, the "pervasive personal monitoring system" may give advice for a more healthy way of life, reliable automated alarms on current or upcoming critical medical conditions, or even automated interventions like prophylactic drug administration to prevent medically critical conditions. The main advantage over current medical monitoring technology will by an integrated approach that is focussing not on the current state or performance of individual organs, but taking into account the complex interactions between many sub-components of the human being as well as the time evolution of all parameters and of their interaction network. This approach of "systems physiology" may also lead to a better theoretical understanding of the complex interactions in human organ systems and their restriction under pathological conditions.

To enable such studies and applications, the science community has to develop techniques that quantify the relations (interactions, causality) between possibly noisy and unreliable signals of very different nature and between parameters characterizing these signals. Interesting information is often hidden in different properties (characteristic parameters) of a signal: averages values, fluctuations, value distribution, auto-correlation properties, spectral power in certain bands, coherence of oscillations, entropy measures, etc. Additional information relevant for the state of the complex system is hidden in the relations between different signals (cross-correlations, cross-modulations, time-delay stability, synchronization, etc.) including directed (causal) and/or time-delayed interactions and control chains or feedback loops. Several approaches based on statistical physics have been shown very fruitful in quantifying such relations. In addition, a well-arranged presentation of the obtained relations is needed to recognize stylized features (characteristic states) and derive predictions on the future evolution of the complex system. Network theory and techniques for describing and studying the time-evolution of complex networks may represent a useful tool for this purpose. Other possible fields of application for the suggested approach include all kinds of complex systems in biology, geoscience, and financing.

Cognitive and Emtotional Adaptation in the Field

Jennifer Healey

Human-Centric Adaptation

Industrial Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

One of the next grand challenges for adaptive pervasive computing will be to may devices that truly understand our mental and emotional situation and try to accommodate us. You probably know that if your boss is on the phone screaming as someone, that this the best time to ask for a raise, similarly if your spouse is in tears after hearing about the death of a loved one, you know that this is not the right time to ask them to take out the trash or remind them of their weight loss goals. Knowing when it is "a good time" to interrupt someone or to make an appropriate is a sensitive issue that requires a high level of social intelligence. We have begun down this path with studies of interruptability, showing that it is possible to know when a person is driving or in a conversation, the next level of understanding will be an assessment of a person's current state of mind. This assessment will likely be based on a combination of long term individual monitoring and in the moment assessment. To understand a person's state, there are several approaches that may be taken and potentially combined. These include sensor based assessment such as the analysis of motion, voice, video, physiological signals or activity and behavior modeling using sensor information as well as behavioral information that could be collected from monitoring entertainment choices, food choices, computer usage, driving style and adherence to planned calendar events. If we can be successful in enabling computers to assess our mental state, we can better rely on them to assist us in managing communication (summarizing e-mail, routing phone calls), planning tasks, balancing our workloads and arranging meetings in our work and social lives. With a model of our mood and mental state, in the moment and throughout our day, our intelligent pervasive system could create an adaptive plan for us. This would include a plan a menu, start our car and then when we get stuck in traffic and start to get frustrated, the system notices and plans a series of alternative strategies to overcome the situation which could include alerting parties that might be waiting for us, finding the right time to reschedule the meeting, calling in to the meeting, or reminding us that we there is no time rush to get to our destination and suggesting an enjoyable entertainment.

Analysis of Life Patterns for Collective Intelligence

Marco Mamei

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
impacting (10) Importance

During the 16th Century the creation and "diffusion" of novel scientific instruments to observe nature (such as telescopes) and new methodologies to analyze the acquired observations (for example differential calculus) led to an unprecedented number of fundamental discoveries that revolutionized science and the lives of ordinary people. Nowadays the increasing adoption of embedded computing devices, smart phones and Web technologies allow us - for the first time in history -- to collect large scale quantitative information about another fundamental realm of nature: the daily life and daily behavior of people. In addition, new methodologies to analyze data of this nature (e.g., pattern analysis and data mining algorithms) allow us to translate this data into knowledge that will expand our understanding of the human condition at the individual and aggregate levels. One of the most interesting applications of this research considers the "quality of life" and the "life style" of European cities as a set of core competencies that could become a ripe business in the future, for example to shape the construction of Chinese metropolies and to contribute the creation of "stress-free societies". The quality of life of European cities is the result of a massive collective intelligence process in which several individual behaviors and activities cooperate together. Extracting the elements of such a collective intelligence is a fundamental research effort. This research consists in developing new methodologies and tools to analyze what is happening in a city and to identify those behaviors positively influencing the life of individuals. The results of this research might have two concrete types of outcome: a) they can supply concrete and quantitative indications to urban planners on how to reshape the city to induce behavioral schemes aimed at improving quality of life; b) they can highlight and make explicit those collective schemes to citizens themselves, and feedbacks could then influence citizens' behaviors and increase the collective intelligence and quality of life. Both these two outcomes fulfill the idea of human-centric pervasive adaptation, allowing individuals to adapt and improve their behavioral schemes on the basis of the extracted information about their life patterns.

Socially-aware Systems

Mirco Musolesi

Human-Centric Adaptation

Academic Contribution

medium (0) Agreement

mid term (>1 year) Duration

impacting (10) Importance

One of the main challenges for the pervasive computing community is to be able to exploit the availability of the exponentially increasing amount of multi-dimensional information about users and their environment (such as location, personal profiles, social network information) in order to build more effective and efficient socially-aware systems. Sources of information include sensors embedded in phones and other mobile devices, but also online social networks such as Facebook and Twitter. In particular, information related to social interactions and emotional states can be derived by means of mobile phones and merged with other context data in order to extract and infer information about the users. The exploitation of personal information and user-generated contents will allow for the design and implementation of new services and systems with people at their centre. This research programme will be focused on different aspects, both theoretical and practical ones. First of all, we need to define new models of human behaviour based on the available data sources. The main challenge is to find meaningful ways of merging these data to infer high-level information about users. The focus should not only be on the single user, but on his/her community, i.e., his/her social network. Therefore, the definition of new models of social interactions from a network point of view is also of paramount importance. The scale of these models can be limited to the representation of physical group interactions, but it can also be extended to interactions in communities and urban environments. From a system design point of view, the key challenges are related to the implementation of efficient and scalable systems supporting these services by leveraging on the recent advances in mobile technologies and distributed computation. Smartphones are the computing platform of the future: computation can be highly decentralised and distributed to users' devices in order to achieve highly scalable solutions. Nevertheless, at the same time, co mputation has also to be centralised especially with respect to the analysis of group dynamics through the merging of multidimensional information potentially coming from several sources including, for example, user-generated contents on social media and sensor data collected by devices carried by other people interacting with the user. This research agenda is truly multidisciplinary involving many areas of computer science (such as mobile systems, distributed systems, data mining and machine learning), but also psychology, sociology and network science. Applications of these new models and systems include marketing, healthcare, behavioural interventions and, more in general, support for studying human behaviour and social dynamics.

Personal Smart Spaces

Nick Taylor

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

The designers of pervasive computing systems have generally adopted a system view of the world in which a distributed infrastructure is approached by users who identify themselves to one of its nodes and then proceed to make use of the available facilities. This system-centric perspective, which perpetuates the traditional view of the user as an actor outside the system, is in marked contrast to Weiser's original concept of calm technology, which located the user at the very centre of the ubiquitous computing world.

Recent developments in smart spaces have tended to maintain this separation of system and user, although this is not always immediately obvious when a user is completely contained in a smart space. It is only when the user leaves such a space that it becomes apparent that they were merely located in it and not properly embedded within it. In order to realise the true essence of the pervasive computing vision it is necessary to start with the user and provide them with a Personal Smart Space (PSS) which is always with them and moves around with them.

A PSS is a personal area network constructed from a variety of networked components ranging from wearable devices to smart dust. This network constitutes a logical domain which affords a limited set of functionalities which are always available to the user. However, as users move through pervasive environments, their personal domains intersect, or come into contact with, other domains through which their base-level functionalities can be extended and enhanced.

PSSs also address the "islands of pervasiveness" problem. Current trends in the design of pervasive systems have concentrated on the provision of isolated smart spaces via a static, albeit often distributed, infrastructure. While this was an understandable approach to the complexity of the challenge when such systems were in their infancy, there is a danger that it will lead to the evolution of islands of pervasiveness separated by voids in which there is no support for pervasiveness whatsoever. The user experience will be all or nothing, with no graceful degradation from the former to the latter. If the current trend continues then, when users meet outside smart environments, there will be no means for them to share or exploit information, services or neighbourhood resources, no matter how powerful or smart their mobile devices may be.

The Personal Smart Space paradigm is being realised by the PERSIST EU F7 project consortium.

Interaction between Objects Born in Digital and Physical Spaces

Peter Bajcsy, Kenton McHenry and Rob Kooper

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

We have been investigating in the past intelligent spaces such as Hazard Aware Spaces (HAS) and Tele-Immersive Spaces for Everybody (TEEVE). These spaces have the potential to bring together measurements from physical worlds with simulations and artistic creations born in digital worlds. One of the long term challenges in the area of Human-Centric Adaptation is the interaction between objects born in physical and digital spaces.

The TEEVE system provides a way for users in distributed geographical locations to collaborate with each other in a shared virtual space. Such systems consist of imaging hardware, networking, displays, human-computer-interface (HCI) technologies, and computational resources to run it all. The functional summary of tele-immersive systems can be described as an environment enabling real-time cloning of three-dimensional (3D) objects in physical spaces, placing them into virtual spaces and then supporting interactions in virtual spaces. The TEEVE technology advances the long term preservation of information, improves education, and enables distributed art in mixed virtual and physical environments. One of the main constraints we have placed on the TEEVE system is the ease with which people can join the environment, this means that the hardware should be easy to acquire, and easy to install and setup.

TEEVE poses a challenge with regards to interactions between objects born in physical and digital spaces since physical laws do not necessarily apply to objects in virtual spaces. For instance, the artistic communities, including dancers and drama performers, have used the ghost like interaction (i.e., objects not colliding but passing through each other in virtual spaces) to create interesting new art performances or to allow audience members to join the performers without any interference with the performance. On the other hand, playing sports in virtual spaces aims at simulating the same interaction as in physical spaces to replicate the same experience. Thus, a real-time human rendering in a virtual space originating from a physical space can touch a basketball born in a digital space and cause the motion of the ball as in a physical space. With a haptic feedback a human could experience some of the sensation they would feel in a physical space. Finally, training in mixed spaces might indicate collision interaction without any physical collision in order to avoid injuries. As an example of this interaction requirement, when wheelchair basketball players use TEEVE spaces for regaining proprioception they need to know about object collisions but just for the purposes of learning how to control their wheelchair and their body and avoiding injuries when playing in physical spaces. With the current trend of immersive technologies becoming more and more pervasive the question remains what laws should govern the interactions in the mixed spaces.

Privacy and Ethics in Pervasively Adaptive Systems

Sarah Spiekermann

Human-Centric Adaptation

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

If we take western philosophy's Idea of Man, which views men as responsible and mature (German: "mündig"), able to act rationally, and capable of defining themselves through moral autonomy and freedom of choice (Kant 1784/1983), we already establish some high level guidance for how ubiquitous computing systems should be built: They should be built as socially aware actors with an inherent respect for the human race. Respect for peoples' ability to judge for themselves and be assisted by machines where needed. Respect for peoples' desire for freedom of choice and be supported by automation and decision support where appropriate. And respect for fundamental human rights, such as privacy, security and safety. A key research are is thus how to build respect for humans' social values into the fabric of machines, to deepen our understanding of 'value sensitive design'. And in doing so to enhance our knowledge on how to respect cultural specificities in such designs.

Peoples' privacy has been a value of particular concern for the ubiquitous computing era. Its dissolution in electronic space is one of the underlying horror visions in myriad science fictions. It is a politically debated subject and it is a research area today and tomorrow for computer scientists, behavioural psychologists and management researchers. Scholars must develop an understanding of how 'privacy by design' can work while not contradicting business models or overwhelming through its complexity. They must elaborate where privacy engineering comes into play in the system development lifecycles. They can improve privacy enhancing technologies (PETs) so that these are more adaptive to context, highly usable as well as educating. And finally, they can contribute to a better understanding of the privacy paradox that finds that people are not protecting their privacy where they claim themselves that they should. This decision making pitfall to act more rationally in a risk context, may be mitigated. Privacy nudging and privacy priming are key words here and a potentially interesting area to study privacy behavior.

VIII — Sociotechnical Systems

Human behavior depends on many factors. Behavior in a situation, is determined by the type of situation, and the way that this situation is defined (cultural norms, social knowledge, or beliefs). Interaction with the environment including other people in this environment is determined by a series of definitions and rules that the culture imprints. These rules make sense in their social context making human behavior contextually significant. The same reaction can be seen as irrational or completely acceptable depending on the context in which it happens. Thus a context for interaction can be seen as a series of rules or ways of behaving in different situations. These rules not only govern individual's behavior but go beyond defining culture in a broader sense. (How does society behave and evolve? How do groups organize, think and act? How do individuals or groups decide in ambiguous situations?. How and under what circumstances do competition and collaboration arise?)

The broad use of Internet and mobile communication technologies has brought up a new set of interaction ways that are affecting the way people communicate and behave. These technologies have radically transformed the social interaction landscape. This transformation is in turn transforming the concepts and definitions that are the pillars of society and social interaction. These technologies are redefining values and rules of social interaction, bringing at the same time new ways of interaction and rules of social behavior. To study these transformations and get a better insight in future social mechanisms we propose to study social transformation through communication in the digital medium.

To study these interactions we will create a location- and time- aware information space, in which users can place and retrieve information using mobile devices or a web interface. Information in this space is linked to a location and a time, allowing users to record history in the making. Through this information cloud users will be able to explore places at past times, trace trends from their appearance time and place, discover places or times where certain feelings are experienced, or witness past events trough all the information left by previous users.

Transformation of society and the self in the digital age

Adrian David Cheok

Sociotechnical Systems

Academic Contribution

medium (0) Agreement

mid term (>1 year) Duration

impacting (10) Importance

In this location- and time-aware, information space users can input information (ex. feelings, messages, recommendations, facts, news, or warnings) using mobile devices or a web interface to a geographical and time location. The information can come in various forms, including (but not limited to) text, images, videos, or sound recordings.

This information cloud will form an augmented reality space that can be navigated using a mobile device (in situ) or using a web interface.
Using the navigation capabilities of this system, users may:

— Scan the world for feelings, news, or facts; examine locations in terms of how people feel or felt (love, anger, hope, frustration, happiness); observe how a trend started, how it spread; observe the circumstances (what, when and where) in which important events happen.

— Obtain summaries or documentaries of what happened in a place at a certain time; review a summary that uses the images, videos, sounds, and comments that people left; or even journey to the future and observe people's plans and wishes.

— Enter the information space deciding what information they want to see, and from when. For example, a user can "walk" in a place and see what people were thinking and doing or what happened in a place two years ago, or s/he can explore a place where some event happened (ex. political movement, celebration, or natural disaster).

— Search by person posting, by place, by time, by concept, by trend, or others.

Advanced data analysis tools will allow spotting trends, observing their movement, their causes, and triggers. This platform, will allow researchers to:

— Explore social and cultural knowledge. What do people believe? And how people act?

— What activities turn into trends? How social and cultural trends appear? and the mechanisms of their propagation. What triggers the transformation from individual custom to group trend?

— Collaborative creation dynamics, competition, collaboration, etc.. What activities promote social collaboration? Study of group dynamics (how groups are formed, how individuals act in groups and how groups interact).

— How groups act differently, from individuals, and from other groups.

— Discovery of the most typical dynamics of interaction (extraction of rules).

— Aid in the study of social interactions in emergency situations (pandemic, natural disaster, terrorism, economic crisis, etc.).

Beyond Pervasive Computing

Albrecht Schmidt

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Mobile communication technologies had a massive effect on people's lives. It changed in a very short time family life and family routines and it transformed for many citizens the way they work. The current transformation that is in full swing is due to awareness technologies (location awareness, context- awareness, based on sensors such as GPS), as they are currently integrated into communication devices and upcoming high quality and multimodal communication and capture. Providing comprehensive awareness and high quality communication media is going to change people's communication needs and it will starkly transform the way we live and work together in the very new future. Within a few years our notion of local (e.g. talking to a person face to face) and remote (e.g. talking via communication device) will blur as will our understanding of private and public. Taken these changes for granted I see great opportunities and major challenges ahead.

Education is one domain where new products and services will have transformational power. People's abilities are largely dependet on the quality and quantity of education and training they received. Such education may be formal (e.g. kindergarten, school, training course, university) or informal (e.g. playing and imitating, looking over the shoulder, training on the job). Upcoming pervasive computing technologies may completely alter the way we teach and learn - increasing quality and quantity of education for citizens. Imagine technologies that provide one to one tutoring for all, such as a pen that leads the way while learning to write or access to a first person view of colleagues who repaired this machine before. The grand research and development challenge is to create the systems (including human users and technologies) in a way that these systems fit people's capabilities needs and desires and essentially to create a learning experience that is perceived valuable and fun.

Personal security and safety will be a further central domain. Personal safety is strongly correlated with people's perception of freedom. What is the difference between walking home from the bus stop at 4pm in the afternoon and at 1am at night? Most people will feel much safer during the afternoon. If they do feel safe to walk through town at midnight they will feel free to do so, if they are afraid, they will not go on their own, their freedom is restricted. This is one example where pervasive technologies will make a significant change to people's perception and realities. I would expect we can create in the near future systems (again made of humans and technologies), that fit our social and ethical rules and would make feel and become safer in urban as well as rural environments.

Generalizing the above domains we can expect a convergence of entertainment, communication, and information. One building block in this is a massive over provisioning of content. Image only a very simple technology: people wear glasses (or another device) that provide what they see and hear to the rest of the world. If we are able to see in real time what other people see, or if we can share what we see, a massive amount of video streams would be available - many more than the world can reasonably consume. However given access to these streams by context meta-data and extracted features many new applications will emerge - and at the same time we will have to rethink ethics and law and in particular our understanding of privacy.

One research challenge in the above is the question of how to further develop the notion of human perception beyond the current sense people have. How can we constantly see multiple perspectives? Can we experience presence, interaction and participation at different physical places at once? Will we be able to create a sense of foresight of events based on community perception? Or will we have means for predicting the likely future?

Overall this calls for a new research approach in pervasive computing centering on the investigation and development of human machines systems that increase the freedom, utilize the power of communities, harvest the ingenuity of a large number of independent developers, and develops technologies that address people's basic needs.

The first generation of ICT aiming towards autonomic systems and their adaptation was driven by the availability of technology to connect literally everything to everything (Connectedness, 1991-2005). Networks of ICT systems emerged, forming communication clouds of miniaturized, cheap, fast, powerful, wirelessly connected, "always on" systems, enabled by the massive availability of miniaturized computing, storage, communication, and embedded systems technologies. Special purpose computing and information appliances, ready to spontaneously communicate with one another, sensor-actuator systems to invert the roles of interaction from human to machine (implicit interaction), and organism like capabilities (self-configuration, self-healing, self-optimizing, self-protecting) characterize the ICT in this generation.

The second generation of autonomic ICT inherited from upcoming context recognition and knowledge processing technologies (Awareness, 2000-2007), e.g. Situation Aware, Self-Aware, Context-Aware, Energy-Aware, etc., reframing autonomy to be based on knowledge, extracted from low level sensor data captured in a particular situation or over long periods of time (The respective "epoch" of research on "context aware" systems was stimulated and fertilized by the PhD work of Anind Dey (2000), introducing the term "context": "...is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and application themselves."). One result out of this course of research are autonomic elements, able to capture context, to build up, represent and carry knowledge, to self-describe, -manage, and -organize with respect to the environment, and to exhibit behaviour grounded on "knowledge based" monitoring, analyzing, planning and executing were proposed, shaping ecologies of ICT systems, built from collective autonomic elements interacting in spontaneous spatial/temporal contexts, based on proximity, priority, privileges, capabilities, interests, offerings, environmental conditions, etc.

Finally, a third generation of autonomic ICT is approaching, building upon connectedness and awareness, and attempting to exploit the (ontological) semantics of systems, services and interactions (i.e. giving meaning to situations and actions). Such systems are often referred to as highly complex, orchestrated, cooperative and coordinated "Ensembles of Digital Artefacts" (FP7 FET). An essential aspect of such an ensemble is its spontaneous configuration towards a complex system, i.e. a "... dynamic network of many agents (which may represent cells, species, individuals, nations) acting in parallel, constantly acting and reacting to what the other agents are doing where the control tends to be highly dispersed and decentralized, and if there is to be any coherent behavior in the system, it has to arise from competition and cooperation among the agents, so that the overall behavior of the system is

Socio-technical Fabric

Alois Ferscha

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

the result of a huge number of decisions made every moment by many individual agents.", Castellani and Haerty (2009). Ensembles of digital artefacts as compounds of huge numbers of possibly heterogeneous entities constitute a future generation ICT to which we refer to as Sociotechnical Fabric, since they pose challenges both in the technological as well as in the societal dimension (already in 2010 we are facing 10^12 - 10^13 "things" or "goods" traded in global markets; 10^9 humans; 10^9 computers on the internet, 10^9 mobile phones, 10^9 cars, 10^8 digital cameras, etc. on planet, megacities with 10^7 citizens, 10^8 users on Facebook, 10^8 videos on YouTube, 10^7 titles on last.fm, etc.)

Mobile interconnecting people with contextually aware augmented reality

Bruce Thomas

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

The challenge is better connect remote people than with a mobile phone employing context aware augmented reality. Web 2.0 technologies have added to people's ability to stay connect with colleagues, friends and family. This is due to the ease of use, pervasive nature of workstation PC's, and the underlying need for the technologies. These forms of technologies do not scale down to smart phone form factors. New technologies need to be investigated to overcome these issues, but in addition take advance of the mobile nature of people. Sensing technologies can provide a vast amount of context data for systems to use. Where, when, how long, and particular tasks can be incorporated into the systems. Because people self form their social networks, what would be considered invasive technologies may be employed. What are these new sensors, and what new infrastructure is required? Augmented reality can provide in-situ information to the user in conjunction with context aware sensing to bring together people, time, place, and activities.

Augmented reality (AR) provides a first person perspective interface allowing the display and interaction of in-situ information in the environment. A simple example is when a user is exploring a location, the following actions are taken when they wish to interact with their friends. Firstly they view the AR information in the location to determine who and where their friends have been at this location. The user's views this information through a head mount display this size and weight of a set of glasses with an AR overlay. Blue cones indicate trails (a path someone has travelled on in the past) by friends who have been there and are currently online. Choosing a friend to share this experience with via Web 2.0 services, they start sharing the experience. This is an example of going beyond a simple phone call and sharing in a more complete manner with their friends.

Future mobile technologies will have to go beyond the mobile phone to connect people. These would be incorporated into personnel artefacts, such as cloths, watches, eyeglasses, and very small portable technologies themselves. Key to making this work is the ability to connect the person's senses with the remote friend's senses. Phones do this with hearing and to a small extend vision. Moving to tactile and better vision would help. Imagine what it would be like to bring a remote friend into the café' with you to share a cup of java.

Pervasive Computing: An Industrial Perspective

Cornel Klein

Sociotechnical Systems

Industrial Contribution
medium (0) Agreement
mid term (>1 year) Duration
impacting (10) Importance

In the last decade, Pervasive Computing has developed into a prominent, vibrant and established research field. Many ideas of the early days - such as gesture based UIs or location aware services - have been adopted by the mainstream IT industry and are available within commercial off-the-shelf products. The more far-reaching vision of a highly decentralized and distributed computing and communication infrastructure pervading the environment has provided research questions for many branches of computer science, creating an eager and lively research community. While there is no doubt that many fundamental research questions still have to be solved - e.g. related to architectural topics like pervasive adaptation, e.g. related to many enabling technologies for pervasive computing - we argue that there is an urgent need to strengthen the collaborative research with other disciplines:

— Software- and Systems Engineering: How are pervasive computing systems architected, designed, developed, deployed and operated in a predictive way? How are non-functional properties such as dependability, robustness, security, scalability addressed? How is the integration of heterogeneous devices, services and technologies from multiple domains addressed in a cost-effective and scalable way ("systems of systems")?

— User centred research: What are the requirements of end-users and other stakeholders for pervasive services and applications? What level of "penetration" with ICT technologies is accepted by various groups of end-users as well as business users? What are proper approaches towards privacy/security? What level is the proper balance between "fully controlled, predictive" and "intelligent" behaviour for particular use-cases and user groups?

— Business centred research: Many pervasive computing scenarios involve collaboration among stakeholders of many domains and roles (e.g. entertainment industry and building control industry; e.g. end-users, service/network operators, owners of the "environment" such as building owners or public municipalities in smart city applications). What are appropriate business models? What are appropriate technical- and non-technical interfaces, standards and processes for interaction between these domains? What kinds of synergies between the different so-far separated technical domains are possible (e.g. use of a converged communication network, use of common "services")? What new applications which will be enabled by technological advancements, what is their business value?

— Societal: How can pervasive computing contribute to address the main challenges of our society, such as demographic change, urbanization, climate change and the need for sustainable development?

Towards Social-ICT Organisms

Franco Zambonelli

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
challenging (15) Importance

The increasing deployment of pervasive computing technologies will lead to the emergence of a dense digital substrate to externalize and enhance our physical and social intelligence, and make it a pervasive aspect of our individual and social lives.

The importance of emerging collective intelligence cannot be denied, as it is the fact that pervasive computing technologies will make collective intelligence so deeply embedded in our activities to make it impossible, in the near future, to distinguish about what aspects of our "intelligence" are to be attributed to us as individuals, to us as member of the world society, or to us as a organs of a continuous and worldwide ICT-Social substrate. My feeling is that the we are close to the so called "technological singularity" that will take the form of a hyper-intelligent, global, collective and adaptive Social-ICT organism.

The processes that will lead to the emergence of such collectively intelligent Social-ICT organisms are already in place, and the most fundamental and critical challenge researchers will have to face in the future is making sure that the sorts of "intelligence" that will arise from them will be useful and not damaging to users and to humanity as a whole. To this end, key directions to unfold the issue include among the others:

— Understanding the trade-offs between the power of top-down (by design) adaptation means and bottom-up (by emergence) ones, also by studying how the two approaches co-exists in modern ICT systems, and possibly contributing in smoothing the tension between the two approaches.

— Understanding the "power of the masses" principle as far as participatory ICT processes are involved. In particular, this implies understanding how and to which extent even very simple collective phenomena and algorithms - when involving billions of entities - can express forms of intelligence much superior than that of more traditional AI techniques.

— Understanding the issue of diversity and of diversity increase in complex systems and in service/data systems, and how diversity of structure and behaviour is currently accommodated in ICT systems. As of now, most studies focuses on a limited number of different classes, which is far from approximating the diversity of existing systems..

— Laying down new foundations for the modelling of large-scale Human-ICT organisms and their adaptive behaviours, also including lessons form applied psychology, sociology, and social anthropology, other than from systemic biology, ecology and complexity science.

— Identifying models and tools by which individual organs of the systems can influence and direct "by design" the emergent adaptive behaviour of the whole system, or at least of substantial parts of it.

Pervasive Software Ecosystems

Gerd Kortuem

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
impacting (10) Importance

The personal computer revolution was initiated by technical advances but would not have been possible without the emergence of a software ecosystem that allowed software developers to create, market, sell and distribute PC software applications. Similarly, we currently witness the rapid growth of an emerging web ecosystem consisting of interdependent platform, application and service providers that through mutual competition and collaboration continue to advance the ecosystem they are part of. Pervasive computing research has made huge strides in enabling technologies, interfaces and solutions but has failed to address or even recognise the importance of software ecosystems for the success and future growths of pervasive computing. A world of highly interlinked pervasive devices, smart objects and smart environments will only emerge if we succeed in unleashing economic and commercial forces that will create a self-sustaining Pervasive Software Ecosystem (PES) that provides a playing field for commercial (and non-commercial) software developers, providers, distributors, vendors and end-users.

A Pervasive Software Ecosystem can be illustrated by tailorable and end-user customisable smart homes: one or more software providers offer smart home software platforms for exposing a smart home's functionality, independent software developers create and market smart home applications built on top of these platforms, software distributors collect and bundle smart home applications, and inhabitants seek out and purchase entertainment, security, comfort and life-style applications for their home and the objects it contains. Similar and more complex scenarios can be envisioned for whole cities, cars or indeed any other environment with embedded, mobile and distributed pervasive technology.

As of now we do not understand how to create the open, recombinant, pervasive software architectures and platforms that would expose software abstractions for enabling pervasive ecosystems. We do not know how to price the value of pervasive services and applications in an open market place. We do not have business models that would allow pervasive software companies to compete by functionality, service level or quality. Finally, we are missing software usage and licence models that would allow software firms to govern and monitor the use and execution of applications in pervasive environments. Addressing these issues requires interdisciplinary, collaborative research in computer science, software engineering, software business management, economics and law. Key research areas include

— Pervasive software architectures
— Pervasive software techniques
— Pervasive software engineering
— Pervasive software business models
— Pervasive software pricing models
— Pervasive software markets
— Pervasive software distribution models

Competitive marketplaces are one of the best ways to foster product and service innovation. Yet pervasive computing is currently in the stage before the PC revolution: highly complex systems, developed and maintained by experts are used to solve highly specific problems. To break out of this mold and to unleash a wave of innovation, pervasive computing research (and researchers) will need to embrace the concerns of independent commercial software providers and address software business and market issues. If we don't, there is a danger that pervasive computing will largely remain an academic discipline while real-world impact is achieved by other communities.

Evidence-Based Sociotechnical Systems

Jeremy Pitt

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

In government, health and transport systems, and other organisations (social, educational or commercial), there is a dramatic tension between policy-makers and policy-subjects. For example, policy makers at European, national, or metropolitan level may set targets (e.g. for carbon emissions, re-cycling, etc.) which require drastic changes in people's behaviour, attitudes, beliefs, values, etc.

There are, perhaps, two extremes of policy formation. These are goal-based policy making: an executive decision is made on the pre-conceived "needs of society", and policy is formulated to achieve that goal; and opinion-based policy making: by some process of consultation, sampling or dim awareness of media pressure, policy is formulated in reaction to the post-conceived "wants of society".

There are many benefits of, but risks in and imperfections of both approaches. However, Ostrom [Ost90] identifies a number of key properties whereby groups which formed institutions to organise and regulate their behaviour did so most successfully. Of these, two were, firstly, that the individuals affected by the rules can participate in the modification of the rules; and, secondly, that the group members (rather than any external agency) had the 'right' to devise and modify their own rules.

Therefore, a judicious interleaving of social networking, multi-agent systems, machine learning and adaptive systems can be used to create a third way, evidence-based sociotechnical systems, which can provide a way of developing stable systems from control of three concurrent processes:
— Defining targets and policies to achieve those targets based on evidence derived from professional social networks;
— Modelling and understanding the impact of policy on behaviour in micropopulations and 'global tribes', across multiple networks with different social, cultural and regulatory norms;
— Modelling and understanding the feedback effects on both policy-making and behaviour at micro- and macro-levels.

Evidence based systems are not based on 'nudge' or 'wisdom of crowds' but are derived from the scientific method applied to policy-making and enactment.

References:
[Ost90] E. Ostrom. Governing The Commons: The Evolution of Institutions for Collective Action. Cambridge University Press, 1990.

Modern Socio-technical 'Oracles' - Predictive Modeling in Socio-technical Systems

Karin Anna Hummel

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
impacting (10) Importance

The human in the loop of pervasive adaptive systems is an important factor as the human is both influenced by the services provided and influences the services at the same time. Complex system models are capable to integrate this feedback loop but the complexity is not easy to capture. In novel large scale systems such as envisioned smart open street parking space management in urban areas targeted, e.g., to reduce CO2 emissions by reducing parking space search times, the feedback loop becomes evident. Due to reduced parking space search times the attractiveness of choosing the car instead of another mode of transport increases which leads to an undesired increase in parking space search times annihilating the first effects. Focusing on the proper design of such feedback loops in pervasive socio-technical systems opens new directions in research and technology development.

Still, one trend observed in today's networked information systems is expected to impact the system model and to lead to technological innovations in the field of pervasive adaptation, that is, the integration of predictive models. Predictive models are widely used to estimate future shopping behavior as well as, more recently, future crime and emergency scenes in crime defense. In the scope of pervasive adaptation, predictive modeling is appropriate as soon as the services pervading everyday life are proactive, i.e., assistive and adaptive based not only on capturing the current state but also on predictions about the future.

Considering at the same time predictive models which lead to pervasive service adaptations and the feedback loop caused by human behavior and tactical decisions is a challenging modeling task. As the predictive nature of systems is well understood in the ICT domain and human behavior in the application domains of the pervasive adaptive systems such as traffic planning, social science, and life sciences, to name just a few, inter-disciplinary approaches are promising new ways to create systems which avoid malicious and unwanted outcomes. At the same time, remaining challenges and research opportunities are the large scale of the self-organizing systems, partial knowledge about real-world behavior such as limited mobility traces available, and overwhelming complexity of influences to be integrated in the model. On micro-scale, approaches such as agent-based simulation might be suitable initial developments to reduce the complexity by modeling single entities and their interactions, while large scale systems require additional technological innovations in terms of meso- and macro-scale models.

Social Dynamics and their Importance for Pervasive Computing

Mads Haahr

Sociotechnical Systems

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

impacting (10) Importance

Humans are building a global communications network set eventually to bring all members of our species within range for communication. The network is growing particularly quickly along the edges where mobile technologies are used. For years, researchers assumed mobile traffic and node movement to be random, but in reality technology is of course used by people whose behaviours are better described by social models. This realization led to social-based routing protocols in which knowledge about social behaviour is used to make better mobile routing decisions.

Such insights constitute a special case of a deeper realization that complex social dynamics govern the use of our entire global communications network and that this knowledge is likely to be useful in a broader context. Many current systems (e.g., Facebook) are concerned with facilitating social interactions and collecting data about them, but pervasive computing systems contain sensors that allow them to infer much more, not merely about user context, but about the social dynamics that govern their use.

I propose the idea of social reasoning as a key feature of any future pervasive computing environment. Social reasoning allows the environment and its applications to capture information about relevant social dynamics and make (likely probabilistic) inferences about them. In future technology stacks, each layer should perform social reasoning as a computation based on captured intra- and inter-node behaviour on a local scale and use it to make inferences about a much larger social behaviour of the system.

A comprehensive social reasoning model for pervasive computing should include a map of the different social dynamics and the social knowledge that can be inferred for each layer in the stack. We know that spatial-temporal node encounters allows inferences related to the social network structure that arises from the encounters and that long-term or short-term tendencies in spatial-temporal node behaviour can be used to infer knowledge related to the environment. For example, a long-term tendency for people to hold their smartphones in landscape mode in a particular location could indicate a photo opportunity, and a sudden large-scale movement of many individuals away from the same location could indicate a crisis.

Taking the long view, perhaps the most compelling application of socially aware pervasive computing is the evaluation and subsequent refinement of sociological models. If we are to survive as a species, we will need to understand ourselves and our own behaviour much better than we currently do.

Supporting Digital Literacy in Pervasive Adaptation

Marc Langheinrich

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Future pervasive systems will be able to constantly tailor our personal devices, our homes and offices, and even public environments to our needs, based on extensive profiling of our activities and interactions. They will guide us to places of personal interest to us; suggest books to read, movies to watch, and products to buy; shield us from unimportant calls and messages; and find interesting people that we might want to meet.

There is an obvious challenge in this personalization regarding the privacy of the collected information: who is to store all this data, for how long, where, and what is it used for? Privacy issues have long been a staple in Ubicomp research, and there is ample research underway in trying to find ways to communicate anonymously, protect user profiles, regulate access based on context, or obfuscate attributes in a statistically neutral way. However, all these efforts often seem in vain, given the willingness of users to give up substantial personal information in exchange for any sort of useful service. Behavioral economists call this "the Economics of Immediate Gratification", where individuals have a tendency to "under-protect themselves against the privacy risks they perceive, and over-provide personal information even when wary of (perceived) risks involved" [Acquisti, 2004]. Providing more privacy controls and more privacy notices may thus not be sufficient to protect privacy, as individuals will most likely find privacy protect not worth the added effort.

In line with today's efforts to bring "Digital Literacy" to citizens [Eshet-Alkalai, 2004], in order to make them aware of both the opportunities and the risks of information technology, we will need to devise means to make people aware of the risks of pervasive adaptation, so that they are in a better position to perceive privacy issues and act upon them. Specifically, we require both technical tools and educational programs that can vividly demonstrate to individuals both the existence and the predictive power of their data shadow, i.e., of the sum of minute little personal facts that they constantly feed into the adaptive systems around them and the inferences that these sum up to. This could be in the form of a personal agent that constantly tries to interpret the sum of facts that it knows about its owner, and that points out whatever a certain range of transmitted values allow the recipient to infer, e.g., her age, income range, or susceptability to a particular product. By spelling out predictions about themselves to individuals, they might be able to not only learn about the predictive power of these seemingly innocent little facts about them, but also alert them to certain mismatches that their data shadow has with their real self, e.g., that they are deemed a credit risk or put into a different age group.

Once people are able to properly understand their own data shadow and its implications, _managing_ this data shadow -- amending, correcting, or even deleting parts of it -- becomes meaningful (and providing such tools would be the logical next step in this research direction).

References:

[Acquisti, 2004] Alessandro Acquisti. Privacy in Electronic Commerce and the Economics of Immediate Gratification. Proceedings of ACM Electronic Commerce Conference (EC 04). New York, NY: ACM Press, 21-29, 2004.

[Eshet-Alkalai, 2004] Yoram Eshet-Alkalai. Digital Literacy: A Conceptual Framework for Survival Skills in the Digital Era. Journal of Educational Multimedia and Hypermedia, Vol 13, 2004

Smart Cities and Ambient Intelligence

Norbert Streitz

Sociotechnical Systems

Academic Contribution

medium (0) Agreement

mid term (>1 year) Duration

challenging (15) Importance

Smart and Humane Cities

Already in 2008, half of the world's population lived in urban areas and about 70% will be city dwellers by 2050, with cities and towns in Asia and Africa are contributing the biggest growth, according to the latest U.N. projections. Thus, the year 2008 was a turning point in human history, as more than half of the world's population lived in cities. When speaking of an Urban Age that we have entered, people predict that the economic prosperity and quality of life will largely depend on the abilities of cities to reach their full potential. There are many issues at a general level spanning a wide spectrum of challenges that are beyond the theme of this book. Being more specific and making the connections to the role of information technology, we have to orient ourselves via sample application scenarios. The theme of Urban Life Management in Smart Cities can be selected as an umbrella scenario. Based on this, we proposed and developed the concept of "The Humane City" as our vision for the City of the Future and the future of Urban Living. Within this context, we can observe a development from real cities via virtual/digital cities to Hybrid Cities and then transforming them into Smart Cities.

We are proposing to combine the general issues and challenges of future cities in an urban age with the motivation to formulate challenges for future research agendas for information and communication technology (ICT), in particular with a focus on ambient intelligence. This resulted in the question: "How can the realization of a Smart City contribute to reducing and potentially even avoiding some of the problems that are faced by today's cities and in the future?" Or in other words: "How can ambient and ubiquitous ICT help to contribute to Urban Life Management?" This can be analyzed and has to be investigated from the following two perspectives:

1. How to manage a person's / a group's life in today's and future cities?
2. How to manage the urban environment of today's and future cities?

While formulating it as two perspectives, it should be clear that they are not independent; but it helps to identify the different user needs depending on who are the users. The users are the citizens living and working in the city; searching, checking, evaluating and then utilizing the services that are offered by the urban environment with respect to the different aspects of life. On the other side, there are people who are organizing and administering the urban infrastructure so that the services are available for citizens and visitors.We have the vision that our cities should be places and environments where people enjoy everyday life and work, have multiple opportunities to exploit their human potential and lead a creative life. Thus, we call them "Humane Cities".

Research Agenda

Ambient Intelligence represents a vision of the (not too far) future where "intelligent" or "smart" environments and systems react in an attentive, adaptive, and active (sometimes proactive) way to the presence and activities of humans and objects in order to provide intelligent/smart services to the inhabitants of these environments. In order to contribute to overcoming the gap between today's situation and the vision of the future as expressed in the goal "Towards the Humane City: Designing Future Urban Interaction and Communication Environments", we developed in the EU project InterLinktwelve research lines (R1- R12) as constituents of a future research agenda. Due to space limitation, only their headings can be listed here:

R1: Rationale for Humane/All- inclusive Cities (users are citizens)

R2: Tangible Interaction and Implicit vs. Explicit Interaction
R3: Hybrid Symmetric Interaction between Real and (multiple) Virtual worlds
R4: Space-Time Dispersed Interfaces
R5: Crowd and Swarm Based Interaction
R6: Spatial and Embodied Smartness (distributed cognitive systems, outside-in robot)
R7: Awareness and Feedback (sensors, physiological, environmental ...)
R8: Emotion Processing(affective computing)
R9: Social Networks and Collective Intelligence
R10: Self-Organization in Socially Aware Ambient Systems
R11: Realization and User Experience of Privacy and Trust
R12: Scaling (as the major horizontal issue)

Addressing and investigating these research agendas items will hopefully help to realize the vision of a smart and humane city.

Social Networks as a Mechanism for Collaborative Computing

Simon Caton

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Social Network Platforms have profoundly changed the way that we communicate and interact today by allowing us to represent, document and explore our relationships digitally. From a relationship encoded within a social network platform we can infer two things: that the relationship probably represents a real world relationship and that there is some level of trust between the peers. The information captured is naturally personal, but rich with details normally unavailable that can be used to foster and facilitate collaborative environments. An understanding and representation of "trust" is, for example, critically important for collaborative environments. The added value of using social networks as a mechanism for collaboration is that they remove certain complexities from domains like Volunteer, Cloud and Grid computing. For example, exchanges are clearly tractable with a (personally) known and "trusted" individual.

Today research projects have only just begun to investigate the adoption of social network structures for different types of collaborative computing. This ranges from community and scientific portals (e.g. PolarGRID, ASPEN), social storage systems (e.g. Friendstore, AmazingStore, omemo, Social Cloud), sharing of networks and infrastructure (e.g. FON) or even things like the distribution of insurance policies amongst social peers (e.g. Friendsurance). Approaches also exist that work in reverse; where a social network can evolves as a result of collaboration(s) both scientifically (e.g. MyExperiement, nanoHub) and through previous interactions over platforms like Amazon's Mechanical Turk.

However, many pressing research challenges, most of which are interdisciplinary, still remain unanswered and are essential for the full exploitation of social networks as a mechanism for facilitating collaboration.

Pressing research questions in this domain include:

— How can such collaboration platforms supported technically and economically?

— How can collaboration agreements between social peers represented, without using formal, invasive, and potentially deconstructive methodologies like SLAs?

— What are the most relevant and effective socio-economic incentive mechanisms that promote collaboration and discourage anti-social as well as strategic behaviour?

— What are the impacts upon participant privacy when using social networks as a mechanism for collaboration? How can privacy be preserved?

— What practices constitute trust amongst participants, and how are relationships affected by such interactions?

— What are the technical, economic and sociological effects of trust transitivity between different communities?

— What is the smallest relationship that can be trusted? How much trust is needed to leverage the social network as a mechanism for collaboration?

Answers to these questions must also consider the constant adaption of the leveraged social (network) fabrics.

A Phase Change in Collective Intelligence: Co-dependence of the Cloud and the Physical/Virtual World

Thomas D.C. Little

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

We are moving towards a world in which most physical objects (people, animals, vehicles, other assets of all scales) are instrumented with sensing and computing devices. Moreover, the objects and their environments are increasingly enabled by ubiquitous network connectivity. This connectivity allows interactions among objects and their state with the now omnipresent computing cloud. We are only just discovering the applications, benefits, and hazards of these trends. This evolution is creating a codependence between the cloud and the physical devices, a new equilibrium for independently controlled and operated devices and ones that are dependent on the cloud. There are many interesting implications of this evolution that form the basis for a diverse research agenda. Specifically:

— What are the future human-computer interfaces enabled by pervasively instrumented objects that obsolete dated interaction models? For example, much richer context-dependent interactions that are based leveraging rich accumulated data about places and individuals that can have persistent state.

— How much intelligence, and computational effort/cost, will be relegated to mobile, object-centric devices? How much will be embedded into the fabric of the network infrastructure? What is the equilibrium between what is required for autonomous operation and what is provided by the collective knowledge base?

— When can we rely on distributed instrumentation to participate as part of safety-critical control systems? At what levels of safety performance can we adopt computer control to achieve its benefits?

— How do we balance individual privacy with collective societal benefits as policies; and how can we realize and ensure compliance with these policies in a massively distributed system with multiple vulnerabilities and failure modes?

— What impact will such a massively connected system have on economic, political, and societal stability? How can it be designed and managed to realize benefits to humanity while minimizing its ills?

— What of the collective intelligence -- accumulated and immediate -- enabled by the cloud and its associated distributed sensors on all objects. If and when will humans become accessories to this one interconnected machine?

In summary, we are in a period of rapid change in the human ability to connect to collective knowledge and control via physical and virtual means and are increasingly co-dependent on these connections to support our daily lives. It is both exciting and a bit frightening to observe these rapid changes as the technology races towards an indeterminate destination with so many unanswered questions.

Computing Is Material for Pervasive Systems

Tom Martin

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

For pervasive systems, computing is material in two ways. First, pervasive systems must intrinsically involve computing. Second, and more subtly, the computing aspects of the system must be treated the same as any other material that affects the feel and behavior of an object. Designers from various fields--such as industrial design, fashion, and interior design--are comfortable with using computing as a tool, for example, editing photos and drawings or creating 3-D models. But with pervasive computing the time has come for designers to consider computing to be part of the physical properties of the object being designed rather than as a tool to use on the object being designed. With pervasive systems, designers will be able to have the look, feel and sound of an object be determined by computation rather than mechanics, allowing for a richer set of interactions between people and the object than is currently possible. Furthermore, products can be truly mass customized, not in the factory but in use, intelligently molding themselves to a person over time. We are already accustomed to objects that mold themselves to us physically as they age: the soles of our shoes wear based upon how we drag our feet, our blue jeans fade to show what we typically carry in our pockets, chairs that we sit upon become indented, and even the keyboards that we type on become polished where our fingers touch them over and over again. These patinas are created strictly from the physical properties of the object. However, when computing is material, products will have "smart patinas," with their wear patterns determined both physically and computationally. These smart patinas will be created in part by the designers, who choose the physical and computational properties of the materials, and in part by their pattern and history of use by the owner. There are enormous implications for design, computing, and economics when objects can sense and adapt to their current and historical contexts. With pervasive systems, three streams of contemporary thought will converge: Design thinking, computational thinking, and long-tail markets. For all this to come to pass will require people skilled at working in the margins between traditional disciplines. It will also require developing proper tools and abstractions that allow designers to focus on form and function of this rich new material rather than its underlying complexities.

Pervasive Learning

Ulrike Lucke

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Knowledge and education are an increasingly important factor in our society. Required competencies go far beyond pure factual knowledge, but include information literacy, media skills, and self-learning competence, to name just a few. Technology can help to obtain these capabilities, especially with regard to temporal, spatial, or personal constraints. Recent evolutionary stages of e-learning, like mobile learning, helped to overcome limitations in time and space. Other approaches, like blended learning, targeted the pedagogical arrangements in computer-mediated educational processes. However, individualization of e-learning is still a challenging field. This concerns many aspects:

— Adaptation of content to pre-knowledge and learning goal
— Adjustment of teachware to situative context of the learner
— Various ways of access according to preferred learning style
— Consideration of differences in language, culture, or personal handicaps
— Support for collaborative processes despite individual variations
— Flexible bridges between various educational tools and platforms
— Seamless integration of learning into everyday living and working
— Persistent transition between physical and virtual artifacts
— ...

While first well-promising results for some former points are already on its way, there is still plenty of research required to fulfill the later goals, in particular.

Future enabling technologies for learning will be developed in a highly interdisciplinary field. Not only are the technical challenges manifold, for instance: user modeling, content management, context awareness, usability, interoperability. Also, there are close relationships to non-technical sciences, like pedagogies, sociology, economics, and law. The need for considering not only a few, but almost all of these facets is characteristic for the whole research field of e-learning. Finally, because of its practical nature, the design of pervasive learning arrangements requires intense cooperation with educational institutions both from academia and from industry.

Towards a community testbed for pervasive computing

Vassilis Kostakos

Sociotechnical Systems

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

I argue that the urban computing research community is in need of a community testbed where ideas can be developed hands-on and tested in a realistic urban setting. Such testbeds have been particularly successful in other domains and disciplines. For instance, the network and distributed systems community has been particularly successful in establishing a number of such testbeds, e.g. PlanetLab (http://www.planet-lab.org) and Emulab (http://www.emulab.net). Similarly, the wireless sensor networking community has developed testbeds such as the Wisebed (http://www.wisebed.eu), while the security community has developed testbeds such as Deter (http://www.isi.edu/deter) for research and development on next-generation cyber security technologies.

The testbed approach is not unique to computer science, since other disciplines have a long tradition of establishing such testbeds. For instance, the physics community has long relied on huge and expensive testbeds in the form of particle accelerators or telescopes. Similarly, meteorology has developed a testbed for easy and accurate testing of weather forecasts (http://www.ral.ucar.edu/jnt), while atmospheric scientists rely on the AMT testbed to systematically and objectively evaluate new aerosol process modules over a wide range of spatial and temporal scales (http://www.pnl.gov/atmospheric/research/aci/amt).

There exists a set of testbed characteristics that are common to most successfully established testbeds across other disciplines. These include open access to the community, supporting documentation, flexibility and configurability. Hence, for a pervasive computing testbed it would be crucial to provide adequate documentation to the various resources, APIs and tools that the users need access to. At the same time liaison researchers, who are acquainted with the testbed, must play a crucial role in facilitating access to the testbed.

Beyond these generic requirements, there exists a set of further requirements that are more concerned with the needs of the pervasive computing community. A crucial difference with traditional testbeds is that in addition to input technologies, output and feedback technologies are crucial to the development of pervasive computing applications. Furthermore, realism is crucial, which is not always the case with traditional testbeds. One way to achieve this would be to integrate the testbed within a real community, e.g. in a city centre. This, however, would come at the cost of access, since it would be very difficult for researchers to remotely deploy and test an application without being physically present on the ground. This is due to high realism and dynamism of a city, which make it very hard to convey remotely.

Arguably, there are many challenges in getting a community testbed off the ground. However, I believe that such a testbed would be of great use to our community in the long run.

Sustainable Pervasive Computing

Vinod Namboodiri

Sociotechnical Systems

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

impacting (10) Importance

Pervasive computing at scale via portable devices has social implications in terms of electronic waste. For example, there are 4.2 billion mobile phones in use globally, with less than 3% typically recycled according to a study by Nokia. Current mobile phones are replaced every 18-24 months, mainly to obtain devices with better performance. One approach to reduce electronic waste is to increase the lifespan of devices involved. One way to realize this approach is through a cloud-based thin-client paradigm where individual consumer devices would have reduced capabilities acting as 'dumb' terminals with most of the processing done at remote servers. This would allow most upgrades to be done at the servers, with little incentive for consumers to upgrade their devices. This approach would require greater reliance on software upgrades than hardware upgrades, a change that should be more sustainable. Research questions to ask would be 1) What are the implications on performance for users as we rely more on thin-clients? 2) What incentives can be provided to a user to keep using their portable devices (like phones) for a longer duration to reduce device replacements? 3) Industry requires customers to replace their devices frequently for profitability; sustainability requires a reduced frequency of device replacement; can these two requirements co-exist? How?

IX — Quality of Life

Virtual Coaches

Daniel P. Siewiorek

Quality of Life

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

challenging (15) Importance

Initial applications in pervasive and ubiquitous computing provided data and information about procedures to mobile users. With the addition of low cost sensors (e.g. accelerometer, light, sound, temperature) and machine learning it is now possible to identify physical activity (e.g. walking, standing, running, driving) and user location. Initial applications such as identifying proper technique (e.g. exercising, lifting boxes, manual wheel chair propulsion) have been demonstrated in research laboratories for the past five years and are starting to be deployed with the public.

The future will see research exploring interaction as a trusted agent. Since trust and privacy are subjective and vary with individuals, Machine Learning will play a key role to customize the agent to individuals as their abilities and interest evolve throughout their lifetime.

Research will engage both "outward" and "inward" trajectories. Outwardly agents will deduce social situations and interactions providing situational appropriate reminders and advice. Inwardly agents will discern physiological stress and eventually understand what we are thinking. Building on recent fMRI discoveries of common spatial patterns among subjects when thinking of the same word, there are numerous projects processing brain signals in an attempt to understand what people are thinking.

Results from outward and inward research will enable the creation of Virtual Coaches that can observe and provide advice just as human coaches do today. The Virtual Coach can grow in capability as we grow. The coach can provide assistance in all aspects of our personal lives. Achieving this goal will require research advances across a broad range of disciplines including social, cognitive, and computer science.

Trends in gas sensing applied to pervasive computing

Jordi Fonollosa

Quality of Life

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
challenging (15) Importance

Pervasive computing systems are based on the Zone-of-Influence (ZoI), which describes the region of the space within a person or an entity can perceive the presence of another entity, or be perceived by others. Since ZoI defines the spatial outreach to which an entity can interact with another entity, pervasive computing systems are limited by the quality and accuracy of the sensors [1].

An enormous effort has been done to develop gas detection techniques to fulfill the rising needs in the progressive industrialization of society through the course of the twentieth century. Hence a wide range of applications are related to gas detection such as environmental and indoor pollution, diagnose diseases like cancer, combustible gas detection, food quality, and others. However, traditional gas sensor techniques show limited performance on odor tracking in plumes and odor source localization, problems that animals like bees and dogs can complete with impressive results. Therefore, pervasive computing algorithms and applications based on gas sensors are very limited compared to other sensor types like distance, movement detection, orientation, light, motion, or temperature.

New techniques of gas sensing that mimic the biological olfaction system are being proposed and aim to outperform traditional chemical instrumentation in specificity, response time, detection limit, coding capacity, time stability, robustness, size, power consumption, and portability. This outstanding performance will overcome the limitations of traditional volatile sensing and will open new applications in pervasive adaptation. Bioinspired techniques will contribute to odor tracking in plumes and odor source localization and, therefore, the ZoI for gas sensing will be well-defined and a new quality of information and interactions between entities and persons will be feasible in the sense of smell and olfactory perception.

[1]: B. Emsenhuber, A. Ferscha, Olfactory Interaction Zones, Video Paper at the Seventh International Conference on Pervasive Computing, Pervasive 2009: Adjunct Proceedings, May 2009.

Digital Ecologies for a Green World

Mihaela Ulieru

Quality of Life

Academic Contribution

medium (0) Agreement

long term (>3 years) Duration

impacting (10) Importance

The new information communication technologies (ICTs) are enabling unprecedented changes to the quantity and quality of information humans exchange with one another and gather from their environments. In what we refer to as 'digital ecology' theory and practice, research aims to understand and advance the interweaving of humans and ICTs to create a world with social, physical, and cyber dimensions enabling a kind of social network in which humans are not just 'consumers' of data and computing applications. Actors in the social network operating within the new digital eco-system are much more: they are producers, 'players,' and 'inputs' whose interactions use the 'invisible hand' of the market as they interact in complex, interdependent global-scale systems in areas such as energy production and use, and neighbourhood, district, city, and regional transport.

Taken together, interconnected grids of communication, electricity, and transport amount to what we call a 'digital ecosystem' whose integrated and reliable operation will undergird development of the next century's energy-efficient and sustainable cities hosting the institutions and technologies of transformed low-carbon economies. Once the digital eco-system is given a name, it is easy to see its ingredients and potential all around us. With advances in sensor and wireless communications technologies, digital ecosystems are poised to connect and even fill existing and newly created applications connecting different environments thus giving rise to many promising solutions to pressing problems. Examples of these environments include energy and communication webs using software applications enabling users to better regulate power consumption and cost, through behavioural change and automated monitoring of home appliance energy use through smart meter networks, in turn connected to time-of-use pricing applications.
Closely associated environments include intelligent transportation systems whose data are fed to smart power grids, enabling them to forecast and meet demand from plug-in hybrid electric vehicles in a city. At the same time as networks of small-scale communicating devices are interacting to form a "skin" that covers physical and virtual spaces, a second major parallel development of ICTs is concentrating computing and storage at optimal locations. Very substantial economies of scale are driving the deployment of cloud computing, where massive-scale datacentres concentrate high densities of computing resources at ideally sited locations to support software applications on demand. The very large scale of these datacentres results in total savings by factors of five to seven times through reductions in expenses for equipment, power, cooling, facilities and operations.

Together, the proliferation of sensors and communicating devices and the emergence of cloud computing represents a major opportunity to develop applications such as smart power grids, intelligent transportation systems, and next-generation communications and collaborations rooted in large-scale cloud computing infrastructures that will provide the foundation for a prosperous post-carbon society. These applications share a common set of features:

— All are driven by the aggregate behavior of communities of human users.
— All contribute to energy efficiency in turn contributing to reduction in greenhouse gas emissions and overall reduction of the ecological footprint of industrial and private activities.
— All have performance and reliability requirements that if unmet have deleterious socioeconomic impacts to the extent that poorly designed applications can create more problems than they solve, ultimately becoming eco-negative.
— All involve control and management mechanisms that mediate the supply and demand of some critical resource.
— All attempt to estimate the supply and demand for the resource using data gathered by a network of sensors.

Successful management strategies and ICT-based tools will enable the development of digital ecosystems which are responsive to and demonstrable beneficial for their users, using reliable management technologies and strategies which promote and enhance the sustainability of industrial and private activities within sustainable low-carbon cities and their surrounding regions.

Based on the premise that management systems that control or regulate the consumption of large-scale resources can be designed using a common set of principles and methodologies that can be adapted to the specific requirements of an application area, we aim at designing digital ecosystems through an interdisciplinary approach that leverages on the synergies across different application areas to exploit common ICT standards and methodologies facilitating interoperability and promoting economies of scale.

Design Challenges for Pervasive Wellbeing

Oscar Mayora

Quality of Life

Academic Contribution
medium (0) Agreement
long term (>3 years) Duration
impacting (10) Importance

A widely accepted definition of Wellbeing states that "Wellbeing is a good condition of existence characterized by health, happiness, and prosperity of individuals in relation to their inner and outer personal spheres". The former definition implies that in order to reach a satisfactorily Wellbeing State it is necessary to simultaneously fulfill the different aspects related to human condition and their impact towards the overall state of being and feeling well. In this sense, research activity focusing on designing future technologies for wellbeing should consider targeting a series of components or personal spheres and their integrations. Such spheres include different dimensions of wellbeing as summarized below:

Physical Sphere - This corresponds to the "corporeal" dimension of wellbeing. The human body and its overall state of wellness (its weight, its shape, its limitations, the amount and quality of activity, its condition against age, etc.).

Cognitive Sphere - Corresponds to the condition of human mind (its activity, the long and short term memory, the intellectual and consciousness states).

Emotional Sphere - This sphere is related to the mental state and motivation, the opinion about oneself, the attitude towards life, the spiritual and affective states.

Health Sphere - Corresponds to the presence or absence of diseases, the impact of treatments, the dependability on others, the access to healthcare.

Relational Sphere - Includes the contact with the outer world, the quality and quantity of social activity, the participation within traditional or new communities, the role in society and its rewards, the family, the work.

Environmental Sphere - Corresponds to the individuals' insertion as part of a bigger complex environment; the pollution and stress exposure, the everyday mobility, the things to do and the resources to achieve it.

To the moment, all the previous wellbeing spheres have been object of study by several research groups in a separate way. They have been analyzed, sensed and monitored through different kinds of pervasive technologies to assess and predict human behaviors and specific conditions related to wellbeing. In the future a more integrated approach has to be taken into consideration for identifying the commonalities and interrelations among them. In this way, future research on wellbeing and in general on wellness informatics field will include pervasive scenarios integrating different aspects of daily living and their concurrent impact on the different spheres of wellbeing.

To achieve this, future systems will need to integrate and analyze multiple sources of data input, beyond traditional context aware computing paradigm by extending contextual dependencies and interrelations simultaneously from multiple potential complementary viewpoints. This fact poses several research challenges to be taken into consideration such as:

i) how the sensed information will be analyzed to simultaneously match multiple wellbeing criteria?

ii) how will the correlations between different wellbeing spheres be established?

iii) how to estimate the overall wellbeing state when some activities under one context (e.g. partying until late night with friends) may affect positively some wellbeing spheres (e.g. relational) and negatively others (e.g. physical, health)?

iv) how new systems need to be designed in order to maximize adoption from users?

v) which new methods/metrics need to be created/considered to evaluate pervasive wellbeing systems?

vi) which privacy and security considerations will need to be taken into account for managing personal data obtained from the different wellbeing spheres?

Adaptive and pervasive assistive technologies for human activity monitoring

Stefano Cagnoni

Quality of Life

Academic Contribution
medium (0) Agreement
mid term (>1 year) Duration
impacting (10) Importance

The scientific community is intensifying research of ICT solutions that enable elderly people to live as much as possible in their homes, autonomously and safely. Ambient Intelligence technologies, used to assist elderly and disabled people, can increase their self-confidence and optimize resources, by monitoring their everyday actions.

Within this context, commonly referred to as Ambient Assisted Living (AAL), to reach the self-confidence autonomous frail elderly people need, to stay in their homes while maintaining an active lifestyle, it is extremely important to detect promptly risky events occurring to them.

Dangerous situations and events can be detected using a number of different sensors - either permanently installed in the environment being monitored or wearable - actuators and processing algorithms, which can be possibly used in different combinations for each detection task to be accomplished.

Pervasiveness is a key requirement for such systems. Even when the area (typically an apartment) to be monitored has limited extension, no part of it can be excluded from surveillance, and no time can be allowed where the system is off. Therefore, permanently-installed sensors must be massively present in the environment to ensure full coverage of the living areas. Typically, such sensors can be cameras, infra-red proximity and motion sensors, temperature and light sensors, etc.

Adaptivity is an even more crucial feature that such systems must exhibit. From a technical point of view, it may allow sensors to switch their operation mode depending on the situation and on the associated detection algorithm, or on the time of day, and so on. Even more importantly, from a behavioral point of view, algorithms supporting the detection tasks must include the capacity of learning. This allows them to adapt to the different behavioral patterns of the people being monitored, which may, in turn, depend on the time of the day or of the year when the events they analyze take place.

According to these specifications, in particular following the need for a fine-grain distributed sensor/processing node network to accomplish the monitoring tasks, from the point of view of the sensors such systems should preferably rely on cheap and simple devices, possibly embedding the capability of performing some simple processing task. From the point of view of the processing nodes, basing such systems on a number of architecturally simple nodes, able to communicate with one another to achieve adaptive self-configuration, appears to be the road to be followed.

Author Index

Keyword Index