Street Computing: Towards an Integrated Open Data API for Cities

Ricky Robinson 1
Markus Rittenbruch 2, 1
Marcus Foth 2, 1
Daniel Filonik 2, 3
Stephen Viller 4, 1

1   National ICT Australia (NICTA), Queensland Research Lab
2   Urban Informatics Research Lab, Queensland University of Technology
3   Ludwig-Maximilian Universität München, Germany
4   University of Queensland

1 Introduction

The physical city is covered with an increasing number of layers of digital information. These rich sets of information are of increasing importance for decision-making processes in councils and local governments and open the possibility to a wide range of future applications and services which are the subject of intensified research efforts. While previous research endeavours have explored the creation of specialised tools to aid decision-making by urban citizens, councils and other stakeholders (Calabrese, Kloeckl & Ratti, 2009; Paulos, Honicky & Hooker, 2009), our vision goes a step further. We are aiming to address the challenge of helping people help themselves by enabling them to create their *own* decision support tools based upon urban data sources, civic sensor networks, and self-disclosed data (for example, from social networking applications). One of the main benefits of this is shifting the monitoring and processing burden from the user to the computational infrastructure.

The integration of information processing technologies into the environment is being investigated from several different perspectives that we discuss further below. However, our review shows that there is a lack of research dealing with the problem of making these information and computational capabilities accessible to everyday users. In order to address these challenges we introduce the notion of *Street Computing* which draws heavily from the existing pool of research. *Street Computing* tries to develop tools and techniques that will help users help themselves. Metaphorically speaking, it is taking computing to the street by giving the general public – rather than researchers and professionals – the power to leverage the available infrastructure and create solutions tailored to their individual needs. It is inspired by Corburn's "Street Science" in that it recognises the imperative of local urban insights for improving scientific inquiry as well as policy and decision-making.

Local involvement delivers crucial information for solving problems in urban communities and increases public trust, leading to a healthier political system (Corburn, 2005). Furthermore, by providing users with the right tools, *Street Computing* enables citizens to take a more active role in the evolution of their cities and thereby reduces the burden placed on government services.

This special issue of the Journal of Urban Technology brings together five articles that are based on presentations given at the Street Computing workshop held on 24 November 2009 in Melbourne in conjunction with the Australian Computer-Human Interaction conference (OZCHI 2009). Our own article introduces the Street Computing vision and explores the potential, challenges and foundations of this research vision. In order to do so, we first look at the currently available sources of information and discuss their link to existing research efforts. Section 2 then introduces the notion of *Street Computing* and our research approach in more detail. Section 3 looks beyond the core concept itself and summarises related work in this field of interest.

1.1   Sources of Information

In the context of our work, we consider three distinct sources of information, data gathered by sensors, information provided by public institutions and social information created by social (mobile) applications. Hardware that enables sensing, computation, communication, and actuation is becoming smaller and cheaper to manufacture. The integration of these capabilities can be observed on various scales, from everyday objects to smart homes and urban environments. Through networking technologies, the individual building blocks are connected to form large systems with ever increasing complexity. Vast amounts of information are accumulated and processed in such systems. At the same time, we can observe a significant trend towards releasing public information. Recent 'Government 2.0' and 'Open Data' initiatives have led to the creation of public data catalogues that contain data such as flood maps, crime statistics and location of facilities such as parks and public toilets. Last, but not least, the ubiquitous availability and increasing prevalence of social media applications provides access to a rich stream of user-generated content. The ubiquity and volume of this information together with the parallel trend towards incorporating location data into web and mobile applications are neatly summarised with the phrase: *The urbanisation of the internet, and the digitisation of the city.*

1.2   Sensing and Ubiquitous Computing

The diffusion of computational capacities into our surroundings is the subject of a number of research areas. Weiser laid the groundwork for *Ubiquitous Computing* in his article "The Computer for the 21st Century" (Weiser, 1991). His vision called for seamless interaction between users and numerous small computers embedded in everyday objects. Weiser argues that such an approach would give users better access to existing and previously unavailable information. For Weiser, the integration of

computers into the human environment is essential in order to allow users to interact with this information in a natural way. Consequently, he sees *Ubiquitous Computing* as a promising solution for the problem of information overload (Weiser, 1991). More recently, researchers have started to point out the disconnect between this envisioned future and the way technology has evolved. While *Ubiquitous Computing* has already become a reality "in the form of densely available computational and communication resources" (Bell & Dourish, 2007), the devices through which information is accessed are "highly present, visible, and branded" (Bell & Dourish, 2007). Nevertheless, a lot of research has been invested in developing technologies and applications that aim at making Weiser's vision a reality. The range of research topics is broad and covers a large scope of application areas. However, many of the projects limit themselves to "small-scale well-defined patches of the built environment such as smart houses or rooms" (Kindberg, Chalmers & Paulos, 2007). In contrast to that, *Urban Computing* is a relatively recent area of research that looks at the impact of ubiquitous information processing at the scale of a city. The focus shifts from integrating computing into everyday objects towards everyday urban settings and lifestyles (Kindberg et al., 2007). Rather than formulating a particular vision for the future, *Urban Computing* aims to explore and understand the implications of emerging technologies on urban landscapes today (Paulos & Jenkins, 2005). Most importantly, it takes into account the complex and dynamic social interactions that take place in cities. Another research field revolving around the role of users is *Participatory Sensing*, which looks at the potential of computational capabilities in the hands of the general public. Using everyday mobile devices, the goal is to "enable public and professional users to gather, analyze and share local knowledge" (Burke et al., 2006). Unlike traditional distributed sensing, *Participatory Sensing* has to consider issues such as ensuring data credibility, protecting privacy, and encouraging participation (Burke et al., 2006). If these challenges are properly addressed, applications in this space hold the potential to engage and empower citizens in new ways (Paulos & Jenkins, 2005).

1.3   Open Data

Recent 'Government 2.0' initiatives have led to the creation of public data repositories such as data.gov (U.S.), data.gov.uk (U.K.), data.gov.au (Australia) on federal government levels, datasf.org (San Francisco), data.london.gov.uk (London) and data.brisbane.qld.gov.au (Brisbane) on municipal levels, and Transport for London (http://www.tfl.gov.uk/) and the San Francisco Metropolitan Transportation Commission (http://511.org) on the level of service providers, to name just a few. Data contained in these repositories ranges from data that is updated infrequently (e.g. public infrastructure) to real-time information (e.g. weather and flood warnings)[1]. While some of this information was previously accessible through different services, the repositories fulfil two key functions. First, they centralise access to information by

---

[1] http://opendata-tools.org/ provides a comprehensive overview of many of these repositories

providing all information through a specific URL. Second, they simplify data access by providing data in a common format and provide a legal framework for the use of data.

Public data is commonly released in a number of formats. These include generic file formats such as CSV, as well as data-specific formats such as DWG and KML for data relating to geographical information systems and maps. In addition to these pre-existing formats, there are a number of protocols specifically designed for open data (e.g. http://www.odata.org/) as well as frameworks that allow public institutions to easily host and publish open data (e.g. http://ckan.org/). Orthogonal to these technical concerns of data storage and dissemination are legal frameworks designed to govern the use and licensing of public sector information. Creative Commons (CC) licenses, such as "CC Attribution 3.0"[2] simplify data access and licensing issues.

1.4   Social Mobile Information

Last, but not least, the ubiquitous availability of social media applications provides access to a stream of user-generated content that includes reviews, feedback, opinions and reporting of incidents, such as traffic incidents, municipal repair requests (e.g. Foth, Schroeter & Anastasiu, 2011) and emergency information. While individual applications serve specific purposes, it is the mining of various social media streams that allows developers to tap into a rich set of interrelated user-generated information.

2   The Street Computing Approach

In the following section we describe our vision for a broad set of research subsumed under the term *Street Computing*. We have begun to implement parts of this vision and are aiming to implement further research projects in this research area with our partners in local government and industry.

We exemplify our research approach by considering the problem of *urban mobility*. We are interested to determine how people currently make decisions concerning their daily travel plans, and how their *in situ* decision-making might be improved by enabling them to construct highly personalised travel planners by combining public transport timetables, real-time traffic updates, ride-sharing information, social network data, personal and office calendar, and any other information deemed relevant by that individual.

In decision-making for daily travel, several common practices can be observed. Commuters use web-based trip planners offered by public transport providers, who may also offer a version that can be accessed on a mobile phone. Some of these providers also offer real-time information on their web sites and bus/train stations.

---

[2] http://creativecommons.org/licenses/by/3.0/au/deed.en

People will often print out timetables to put on their walls for reference. Furthermore, they may tune into traffic reports on the radio to keep up with traffic incidents which may influence their choice of transport mode. Recently, some public transport providers have begun to use services such as Twitter to make announcements about cancelled services, track work on railways and so on. Twitter has also been used for "crowd-sourcing" traffic incidents.[3] Finally, personal events recorded in calendars or diaries must also be consulted, often for the purpose of putting hard bounds on arrival time at the destination. These non-provider sources of information play a large part in decision-making, and illustrate that the travel planning process is different for each person.

One limitation with current practices is the necessity for commuters to monitor all of these various information sources and combine them in their head. A second limitation is that existing technology solutions to the travel-planning problem take a provider-centric view rather than a commuter-centric one, meaning that commuters are offered trip planning tools that are not tailored to their individual needs and circumstances. Note that merely having access to real-time information about bus and train arrival times does not solve this problem.

The state-of-the-art in trip planning is the Personal Travel Assistant (CUD, 2009) from Cisco and its partners in the *Connected Urban Development* program, which does endeavour to provide a personalised decision support tool. However, it is limited in that users cannot add their own information sources (such as calendar information) to the tool. Furthermore, the solution is specific to urban travel, rather than a general platform that can underpin solutions to many problems in the urban realm.

We are interested to explore end-user programming techniques to enable users to construct their own decision support tools. However, our solution is general, and will have many applications beyond the problem of urban mobility that is being considered here.

There are three major elements to our approach:

(1) interaction techniques for end-user composition;
(2) a framework for composition; and
(3) the urban computing substrate.

2.1   Interaction Techniques for End-User Composition

There are numerous examples of end-user programming in modern everyday life, ranging from simple but inflexible approaches to more sophisticated techniques that are still beyond the grasp of the lay user. We outline some of these approaches here, moving from the simple to the more complex:

---

[3] See, for example, SAP Research's iPhone app AUS Traffic, and the #bnetraffic hashtag on Twitter

- Setting the time or recording schedule on a VCR or PVR (personal video recorder). The task has been simplified with the advent of G-codes – a short numerical code representing a TV program that can be typed into the VCR or PVR;

- Programming-by-example in desktop applications such as the Microsoft Office suite, whereby the user "records" a sequence of actions so that this sequence can be invoked in a single step at a later time. A generated script in a high-level language is often used to represent the sequence;

- Repetitive tasks in desktop environments can be automated with the use of tools such as Apple OS X Automator, in which users compose a workflow from a set of actions that can be linked together. In the same spirit, but for a different purpose, we are beginning to see the creation of tools to simplify the task of creating "mash-ups" from web-based data. Tools such as Intel MashMaker[4] and Yahoo Pipes[5] fall into this category.

While end-user programming has been shown to be feasible for customising personal spaces (Chin, Callaghan & Clarke, 2006a), creating video/audio capture and access tools for the home (Truong, Huang & Abowd, 2004) and configuring home multimedia systems (Newman, Elliott & Smith, 2008), our approach is to create an experience resembling something that most people are already familiar with: meal preparation. Just as people go shopping for ingredients that they bring home to combine in a particular way to create a meal for themselves, our approach involves visually exposing data and computational elements in the physical world, which can be "gathered" by ordinary people, and then combined by them in a specific way to achieve something that is personally or socially meaningful.

In contrast to existing end-user approaches to combining data and processing, our approach, which we term *conceptual composition*, is:

1. *Situated* – users carry with them the context within which a computational element was captured, which serves as an aid to the user at the time of composition;

2. *Physical* – we want to provide users with an experience similar to that of meal preparation, which deals in physical objects;

3. *User-centred* – our approach considers the manner in which users approach similar tasks. For example, trial and error is one strategy commonly deployed by people in trying to achieve a particular end result from a combination of elements (in cooking, for instance), and so it follows that we should investigate how our

---

[4] http://mashmaker.intel.com/
[5] http://pipes.yahoo.com/

solution can support such behaviour. Other less obvious constraints and behaviours will be discovered via a set of user studies that underpin our approach.

In the urban mobility use case, this approach might be realised by embedding symbolic representations of components (such as bus timetables and real-time traffic information) in the environment, and to enable users to "capture" the elements on their mobile devices for later use in a composition. Initially, these symbols might be unique identifiers such as QR codes. Later, and depending upon our findings from user studies, we may evolve this capture technique to use object recognition in combination with augmented reality, so that components can be identified on the screen of a mobile phone as one walks around, and then "dragged" into a composition or onto a clipboard for later use.

2.2   A Framework for Composition

The proposed framework integrates tightly with both the interaction techniques for end-user composition, and the computing substrate below it.

The composition framework maps the user specified composition to an executable representation and oversees the (distributed) execution of the composition. However, the framework will also support directly programmed applications, which might be developed by a skilled engineer as opposed to an end-user. In this sense it is analogous to a web development framework such as Django[6] or Ruby on Rails[7], or a toolkit for context-aware computing (Dey, Abowd & Salber, 2001; Henricksen & Indulska, 2006).

Unlike web development frameworks, an underlying network of dynamic data sources (that is, the urban computing substrate, described below) rather than relational databases will back the applications. Furthermore, the processing and rendering components of a composition may be executed locally or remotely according to need and environmental constraints, for example, does the local environment have an appropriate device or public display for rendering visuals?

We will explore several different techniques for mapping user-specified compositions into executable code. One possibility is that the graphical or tangible objects exposed to end-users are tied directly to corresponding software objects within the framework, and workflow creation then becomes a matter of traditional object-oriented composition. Another approach is to use meta-programming, in which the end-user compositions generate program *closures* within the framework, which can be executed by the local framework instance, or sent to one or more remote sites (i.e., in the cloud). We will see that the latter approach is consistent with the proposed architecture of the urban substrate. Because one or more remote machines may

---

[6] http://www.djangoproject.com/
[7] http://rubyonrails.org/

execute the composition, the generated code should be expression-limited so that safety in relation to resource consumption and so on can be proved simply.

Motivated by the urban mobility use case, we use the framework to explore the concept of "smart" compositions through a significant places learner (Hightower, Consolvo, Lamarca, Smith & Hughes, 2005; Kim, Hightower, Govindan & Estrin, 2009; Lane, Lu, Eisenman & Campbell, 2008; Zhou, Frankowski, Ludford, Shekhar & Terveen, 2007), and an extension called *significant situations*: re-occurring situations that hold subjective meaning for an individual or group. This investigation of significant situations is more difficult than the creation of the framework itself in terms of implementation and the range of second order challenges that flow from it.

We assume that significant situations may be predicated upon arbitrary components or higher order compositions, selected or created by the end user. We will explore a number of approaches to situation learning and recognition. Brdiczka, Crowley and Reignier (2007) employ Hidden Markov Models (HMM) to *learn* situations, which are updated with user feedback. However, their work does not relate to end-user programming or decision-support tools; rather it relates specifically to context-aware service selection, whereby particular services are invoked in particular situations. Other candidates include Bayesian networks, Dempster-Shafer Theory and Support Vector Machines, some of which have been employed previously in the closely related task of context reasoning (Bettini et al., 2010).

Leveraging the early work of Lane et al. (2008), we also want to investigate the potential to share data and models between people within the same social network to improve the performance of our recognition algorithms. To this end we will attempt to derive an ontological basis upon which to formulate an understanding of situation models. In other words, is it possible to determine whether two or more situation models, independently specified (via conceptual composition) and trained by different people, are equivalent, or related in some other way (subsumption, for instance), such that we can decide whether or not to share models and/or training data amongst users?

Given a formal basis for the sharing of models, we are obliged to consider the privacy implications for such sharing. In practice, there will be a trade-off between improved user experience through superior situation recognition and privacy. A related question is under what circumstances an inference can be disclosed to other users, that is, if a situation is recognised, who must/should/must not be notified? Drawing upon existing work in privacy control (for example Boyle & Greenberg, 2005; Palen & Dourish, 2003; Wishart, Henricksen & Indulska, 2005), we will investigate how obfuscation can be applied to the underlying data or the user interface to reduce the resolution of situation recognition. For instance, an ontological model of significant situations will enable the disclosure of coarser grained situation information in the case where the finer grained situation (say, "waiting at the bus stop on Moggill Road") is subsumed by a coarser grained one ("somewhere in suburb X").

2.3 The Urban Computing Substrate

The urban computing substrate is the collection of computational and data elements together with the protocols that underpin the above. Current architectures for distributed data-driven applications generally follow one of two approaches (Henricksen & Robinson, 2006). Middleware for sensor networks take a bottom-up approach that starts with the capabilities and constraints of the hardware platform and endeavours to provide software engineering abstractions that assist with extracting data from the network without requiring application developers to deal with low-level hardware and networking issues. An alternative is the top-down approach, in which a deep understanding of application requirements is a primary driver for the design of the middleware. To some extent, this is the approach that has been taken in the field of *context-aware computing*.

What is now required is a converged approach, which lends itself to the sophisticated high-level programming and query abstractions typical of context-aware computing environments, as well as the distributed execution environments offered by middleware for sensor networks. Coupled with the framework described above, our substrate will provide the execution environment and messaging primitives to realise the conceptual compositions defined by end-users and developers. To this end, we propose an architecture that is similar in style to Computational REST (Erenkrantz, Gorlick, Suryanarayana & Taylor, 2007), or CREST, which advocates computation rather than content as the fundamental unit of transfer.

However, CREST is silent on publish/subscribe models of communication, which are prevalent in the literature on middleware for sensor networks and context-aware computing due to the loose coupling they afford (that is, subscribers need not have any knowledge of the identity and location of publishers), and which are desirable in real-time decision support systems. We therefore explore asynchronous extensions to CREST, to enable scalable, selective, real-time updates from streaming data sources or lower order compositions. In particular, we investigate computational subscriptions and notifications. That is, in keeping with CREST philosophy, subscriptions and notifications will be executable code (closures), meaning, for example, that publishers can export the processing overhead due to data conversion from one format to another to the subscriber (a generally more scalable approach) or any intermediary host, and that the subscriber need not be pre-programmed with the particular data conversion algorithm. A further use for this architectural style is that notifications may simply be pointers to other computations or data instead of containing the data itself, and the client application need not have any knowledge that this is the case.

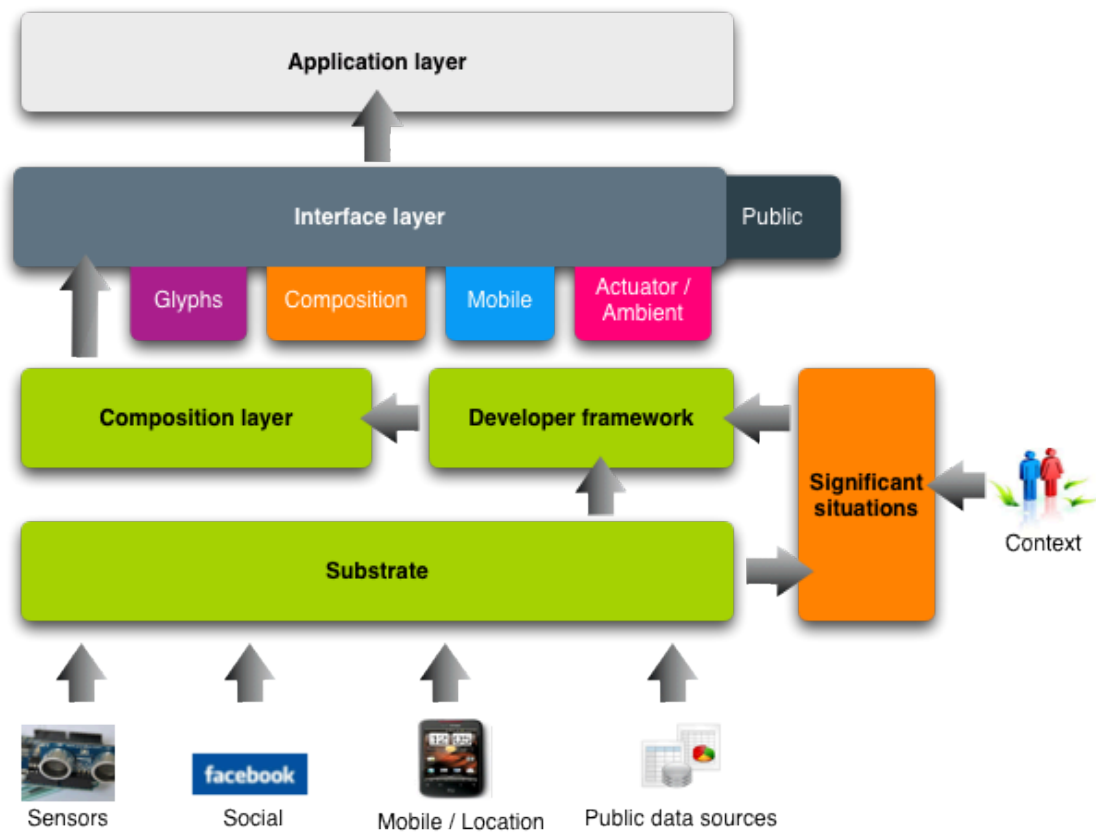Figure 1 depicts the different components of our research approach.

Figure 1: Integrated Open Data API - Components

2.4 Individualised Data Stream Composition and Visualisation

We have started to implement some aspects of the larger research vision of Street Computing. We have developed an iPad application that supports individualised data stream compositions and visualisations, in form of a dashboard, that allows users to easily combine and display different types of data sources. The application is aimed at individual home users and allows access to data sources such as home energy data (e.g. through Pachube) and social media data. The project acknowledges that individual users have different preferences with regard to the style in which they want to present certain data (e.g. graph vs. smiley face animation for energy data) and with regard to which data they want to display at the same time. While the dashboard does not yet support full data integration, it is a first step to explore the individualisation of data streams. Figure 2 displays some of the current mockups.

3 Related Works

The following section presents the major challenges and open research questions in the area of *Street Computing*. These issues were identified by reviewing literature from related research areas and analysing projects that are positioned in this particular space. The review aims to point out reoccurring problems and group them so they can be addressed in future research.

3.1 Integrating Heterogeneous Data Sources

Already today, our urban environments are monitored and analysed through distributed computing and communication technologies. Complex and mostly independent systems support our everyday activities in the city, such as public transportation, mobile communication and financial services. Further, various corporate and government entities collect data about traffic, pollution or weather. It is foreseeable that this computing environment will incorporate an even wider variety of devices and services from different manufacturers and developers. In order to accomplish the goal of making this wealth of information from urban systems accessible to users, *Street Computing* needs to address the issue of working with heterogeneous data sources. Therefore, to provide a solid foundation for future projects it is crucial to achieve platform and vendor independence, as well as architecture openness (Lee et al., 2007).

Zambonelli envisions those separate urban systems as individual organs that can "collaboratively contribute capabilities and services aimed to support increasing quality of life levels" (Zambonelli, 2011). He advocates the development of proper collaboration and middleware tools that allow "all the components of the infrastructure (…) dynamically and adaptively connect and collaborate with each other" (Zambonelli, 2011: 578). Such tools would greatly simplify the development of applications for *Street Computing*. Currently, most projects resort to creating this functionality from scratch, investing considerable time and effort. Providing adequate solutions to this challenge would benefit scientific progress by allowing researchers to focus on other questions that lie beyond.

11

For instance, the *Weave* architecture is one specific effort to develop standard APIs for services offered by sensing fabrics. Its goal is to simplify the design of applications across heterogeneous sensor networks that are independently managed. It recognises that it is often necessary for applications in urban sensing to "fuse information from sensing fabrics extant in the environment, to meet requirements for which the fabrics were not designed a priori" (Kulathumani, Sridharan, Ramnath & Arora, 2008). The work distinguishes between a generic API for urban sensing and vertical APIs specific to application domains. The architecture is put to the test by focusing on the domain of search applications and applying it to concrete scenarios (Kulathumani et al., 2008).

While it may not be feasible or even desirable to try and formulate an all-encompassing generic standard for *Urban Computing* – given the wide variety of application domains and problems – it would be helpful to have guidelines and best practices that manufacturers and developers can follow. Certain aspects can certainly be abstracted, such as the protocols for publishing and subscribing to data streams. However, even with a successful attempt at standardization, it is still necessary to have ways of integrating legacy and non-conformant systems.

3.2   Extracting Information and Knowledge

As the technologies for gathering data are widely deployed within urban areas, including smart homes, the leading question becomes how to process the collected data. The relationship between data, information, and knowledge is explored in information systems and knowledge management literature. It is often represented as a hierarchy in which the higher levels are increasingly rich with meaning. While this concept is generally agreed upon, there is less consensus about "the processes that transform elements lower in the hierarchy into those above them" (Rowley, 2007). According to the classical definition by Ackoff, knowledge allows us to apply data and information to give answers to questions of 'how' (Ackoff, 1989). Such questions lie at the core of the problems that urban communities face, for example "How to reduce traffic congestion?", "How to support healthy and sustainable lifestyles?", or "How to plan and develop neighbourhoods with a high quality of life?".

Comprehensive ontologies suitable for *Urban Computing* are a critical prerequisite for applying semantic technologies. In turn, these technologies hold enormous potential for *Street Computing* by semantically interpreting the sensor output and thereby allowing non-technical users to benefit from data which would otherwise be meaningless to them. In the design of the search API for the *Weave* architecture proposed by Kulathumani et al. sensor fabrics are viewed as providers of an extensible database of queryable objects. As part of their future work, the authors aim to extend the search functionality to enable queries with richer semantics. In order to accomplish this goal, they are planning to develop "a richer ontology for urban sensing that builds up from physical phenomena, sensing signatures, detectors, and richer classes of objects" (Kulathumani et al., 2008).

Effective decision-support systems require actionable information and knowledge about the situation in question. In order to provide this information, algorithms need to be developed that can combine static knowledge about the urban environment with the dynamic data streams originating from sensors. Della Valle et al. propose a system that adapts reasoning techniques to work with the resulting time-varying knowledge. Their solution combines semantic reasoners with throughput-efficient data stream management systems (Della Valle, Ceri, Barbieri, Braga & Campi, 2008).

The *Semantic Streams* framework by Whitehouse et al. presents another approach that supports the quick extraction of semantic information from sensor output. The framework is based on a declarative language for specifying and composing inference units, which are "minimal units of sensor data interpretation" (Whitehouse, Zhao & Liu, 2006: 5). Existing inference units can be dynamically composed to generate new interpretations of sensor data. This allows non-technical users to pose queries over semantic interpretations of sensor data without having to write code performing the inference themselves (Whitehouse et al., 2006).

3.3   Supporting Mashups and End-User Programming

An area that is still largely unexplored in the context of urban computing is the potential for end-user programming. This aspect is a fundamental part of *Street Computing* and its vision of empowering users. Making it possible for the general public to leverage the existing technical infrastructure and create solutions tailored to their individual needs requires the development of new tools and techniques. A user-centric approach is necessary to make the available functionality easy to understand and use.

The issue of end-user programming can be addressed from several different directions. There is substantial room for improvement when it comes to querying the sensing environment. Users need better ways of retrieving the information relevant to their individual problems. Query languages based on semantic interpretations of sensor data permit a more natural way of inquiry by allowing the use of high-level concepts. This relieves the burden of making low-level decisions based on the raw sensor data from the user. Such an approach has the additional benefit of introducing an extra degree of freedom to the system. The user does not specify which exact sensor data is used and how it is processed to infer the desired information (Whitehouse et al., 2006).

Another means of approaching end-user programming is through the composition of existing data and services. This has recently come to the centre of research in the realm of web services, where such compositions are referred to as mash-ups. Several different paradigms of end-user programming have been explored with regards to creating mash-ups. Among those are simplified domain-specific languages, programming-by-example, and visual programming based on the data flow metaphor. The latter is a rather prominent paradigm that has been employed in several

commercial implementations, such as *Microsoft Popfly* or *Yahoo! Pipes*. In this approach, the composition is represented as a graph of components, each of which has input and output parameters. A component transforms the data received in the input parameters and stores the result in the output parameters. By creating connections, the output parameters of one component act as input parameters of another, resulting in a design pattern similar to the pipeline or pipe and filter technique popularised by Unix command-line shells (Biörnstad & Pautasso, 2007; Wong & Hong, 2007).

The *SensorMasher* platform is one specific example for adapting the concept of mash-ups to sensing environments. It specifies how to publish sensors on the web and access their associated data through URIs. Further, it provides ways of linking these published entities to ontological concepts. This information is used in the visual composer to provide an intuitive interface for navigating and exploring sensor data sources by following semantic links. Once the user finds the relevant sensor data sources, he can use the work flow editor to visually combine the data and apply data processing operators. The resulting mash-up can be published and put to work in subsequent compositions (Le-Phuoc & Hauswirth, 2009).

Further research is necessary to explore ways for users to control the ubiquitous computational capabilities that form the basis for *Street Computing*. Chin et al. extend programming-by-example to a ubiquitous computing context by allowing users to "show" the system the desired behavior "via natural physical interaction with the environment" (Chin, Callaghan & Clarke, 2006b: 325). The system deduces rules based on the actions that users perform with the physical devices themselves or with the corresponding interface elements. Beyond that, other approaches incorporate tangible user interfaces to represent key programming concepts in the course of the interaction between the user and his surroundings (Chin et al., 2006b; Hague, Robinson & Blackwell, 2003).

3.4   Interfaces Supporting Composition

One of the challenges of developing interactive technologies for the urban environment is the diversity of interaction technologies, modalities and settings in which the interaction takes place. In order to support composition of programming in-the-wild we need to understand both interaction mechanisms and available technologies as well as the specific use context people operate in.

A large body of research in Ubiquitous Computing and Human-Computer Interaction has laid the groundwork for understanding how computing can manifest in a post-desktop world. We will focus on three major aspects: mobile interfaces, ambient and situated interaction and novel interaction techniques such as gesture-based and organic user interfaces.

Post-desktop and urban interfaces

The now near-pervasive availability of mobile devices[8] has opened the door to new types of interactions and services. The notion of location and location-based services lie at the very heart of ubiquitous computing research. Researchers in this field have explored services ranging from electronic city guides, (e.g. Cheverst, Davies, Mitchell, Friday & Efstratiou, 2000; Paay, Kjeldskov, Howard & Dave, 2009), education (e.g. Underwood, Smith, Luckin & Fitzpatrick, 2008), location-based games (e.g. Benford et al., 2004), to location-based notification systems for police officers (e.g. Streefkerk, van Esch-Bussemakers & Neerincx, 2008). There is a growing number of commercially available social mobile applications, which use location data to maintain social networks (e.g. Foursquare, Google Latitude, Brightkite, Plazes, Fire Eagle, Zkout, Rummble) and predict social trends (e.g. CitySense[9]).

An important development in this context is the use of augmented reality in combination with mobile devices in order to overlay to city with an additional interactive information layer. For instance, iCam (Patel, Rekimoto & Abowd, 2006), a location and orientation aware device has been used to demonstrate the annotation of real-world objects (without the necessity of static labels). Metro Paris Subway[10] is an iPhone application that overlays real-world environments with additional information enabling service discovery and way showing (see Figure 3).

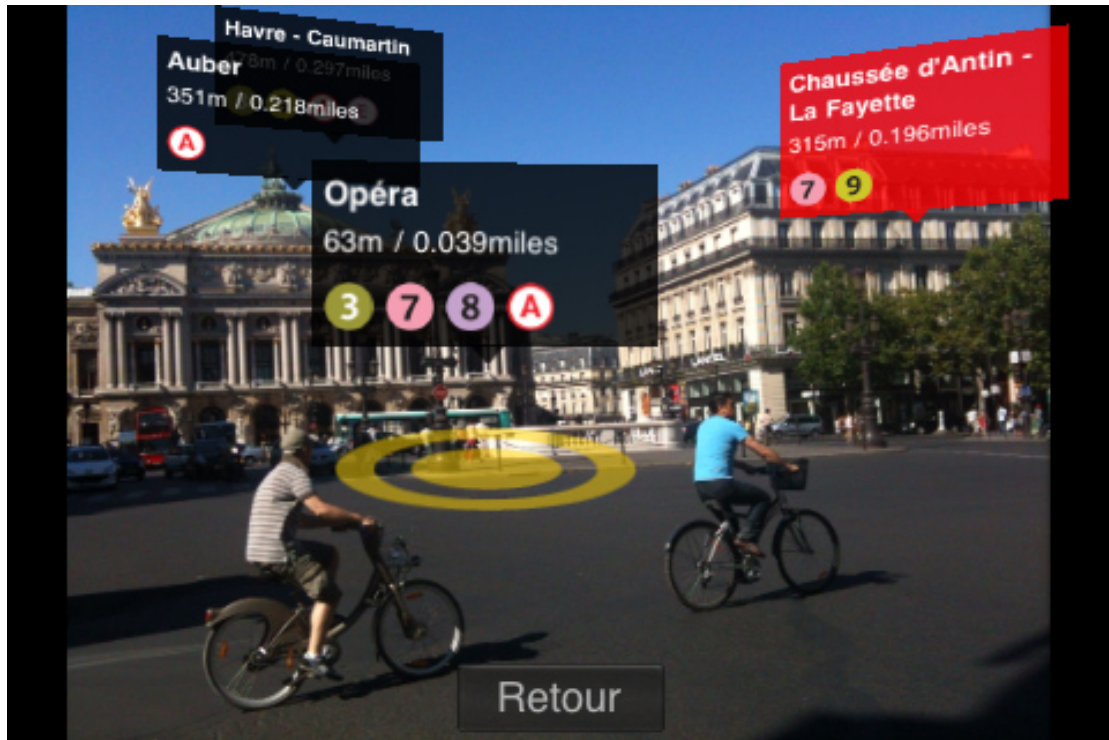[8] The Australian mobile penetration rate is predicted to reach %122 by 2012 (ITWire, http://www.itwire.com/content/view/18424/1231/)
[9] http://www.citysense.com/
[10] http://www.metroparisiphone.com/

Figure 3: Metro Paris Subway[11]

Figure 4: EyeStop by the MIT Senseable City Lab[12]

While personal mobile devices have been the predominant platform for interaction away from the desktop in recent years, ambient and situated displays play a similarly important role in Mark Weiser's vision of ubiquitous computing. Advances in interaction technology, like the availability of large-scale multi-touch wall displays have led to an increased deployment of interactive displays in urban environments.

[11] http://www.metroparisiphone.com/images/photos/9.jpg
[12] http://senseable.mit.edu/eyestop/01.jpg

CityWall is a large multi-touch interactive display installed in Helsinki, which displays information about events happening in the city. Peltonen et al. (2008) studied how large interactive displays get shared between multiple people and addressed issues such as crowding, teamwork and conflict management, as well as the questions of how content from personal mobile devices get shared on public displays (Peltonen et al., 2007). The MIT Senseable City Lab has designed an interactive bus stop that is targeting a large-scale rollout in Florence, Italy (see Figure 4). These EyeStop[13] stations will allow users to interact with bus locations, timetables, bus proximity, community boards, route plotting functions and advertisements. In addition, there is an increasing amount of research into non-standard urban interfaces such as, e.g., urban pixels (Seitinger, Perry & Mitchell, 2009) or façade interaction (Dalsgaard & Halskov, 2010).

While each of the services and technologies discussed here represent interesting approaches to interacting with information away from the desktop, they focus on isolated modes of interaction. The EyeStop concept goes further in that it combines different display technologies (e.g. E-ink and LED displays) with different modes of interaction (web-browsing, timetables, ambient display, etc.). Our approach differs from the above in that we are seeking to develop a platform that grants users access to the underlying layer of information and allows them to not only compose their own application, but also to use existing public input and output mechanisms available in the city environment.

Interaction techniques for composability

Tangible computing has been shown to be successful in a wide range of situations, including early learning, e.g., *Token Bowl* (Martinussen, Knutsen & Arnall, 2007), *jabberstamp* (Raffle, Vaucelle, Wang & Ishii, 2007) and *FlowBlocks* (Zuckerman, Grotzer & Leahy, 2006); editing media content, e.g., *Mediablocks* (Ullmer & Ishii, 1999); disaster planning, e.g., the *Tangible Disaster Simulation System* (Ben-Joseph et al., 2001); and urban planning, e.g., *Urp* (Underkoffler & Ishii, 1999) and *Illuminating Clay* (Piper et al, 2002). Tangible User Interfaces (TUIs) couple physical representation with digital representation (Ullmer & Ishii, 2000). TUIs are particularly suited to the tasks of composition as physical objects are used as handles for virtual objects, e.g., *Bricks* (Fitzmaurice et al., 1995). Virtual objects are modified by manipulating corresponding physical objects (moving, shacking, orientating, exerting pressure, etc.). Composition of physical objects is a common task in TUIs and has been explored in many contexts, for instance musical composition (Jordà et al., 2007), and building simple programs through sensor network user interfaces (Merril, 2007).

With regard to composability we argue that the combination of tangible computing with surface or touch-based computing can provide an intuitive, natural interaction

---

[13] http://senseable.mit.edu/eyestop/

between the computer and the human during the application development phase. The *Reactables* (Figure 5) are an example for such an approach as they combine tangible objects with touch-based interaction (multi-touch) on a tabletop surface (Jordà et al., 2007). However, unlike the reactable where the display is mainly used to represent visual feedback on the immediate interaction, the Street Computing approach requires the visualisation of a wider range of information related to specific sensors, services including maps, environmental data, contextualised sensor data and computational elements such as data filters and aggregators. We envisage that this will allow users and developers to explore and simulate the implication of different combinations as part of the process of creating compositions.

Figure 5: Reactable[14]

3.5   Contributing and Sharing Resources

Street Computing aims to develop systems that allow the general public to make use of distributed computational resources to gather information about urban environments. At the same time, the public should be encouraged to contribute resources at their disposal. With increased availability of resources, the potential applications of the system become broader. However, resources may be exclusive or constrained, making it necessary to coordinate access and resolve conflicts. For example, if a user chooses to supply their mobile device as a sensor, battery life would be a resource limiting how the sensing is carried out.

In order to design a system that relies on collaboration and sharing, it is essential to understand the factors that motivate people to participate. For certain users, the driving force might be seeing how their data contributes to a larger whole. In other cases it may be necessary to identify and design novel interaction mechanisms to provide incentives for participation (Cuff, Hansen & Kang, 2008; Zambonelli, 2011). Taking resource constraints into account, one viable approach might be to grant preferred access to those users who also contribute to the system.

---

[14] http://www.reactable.com/images/live2_design1.jpg

The problem of limited shared resources in urban sensing systems has been addressed in different research projects. It becomes especially relevant if the systems provide end-user programming capabilities, which are inherent in *Street Computing*. The *CitySense* project is one such example. In order to tackle this issue, resources are exposed through a custom API, which serves the purpose of restricting the programming model and allowing the underlying framework to perform optimizations such as caching (Murty et al., 2008).

Another approach for distributing load and resolving resource conflicts is taken by Whitehouse et al. and their *Semantic Streams* framework. The declarative specification language provides the capability to define quality of service trade-offs. Their framework uses this information together with its semantic inference capabilities to dynamically select the best available resources for answering a particular query (Whitehouse et al., 2006).

3.6   Providing Security and Privacy

Sharing and collaboration are fundamental components of *Street Computing*. Its applications can only thrive in an environment where users are encouraged to participate and feel comfortable to exchange data. Individual contributions of local and domain-specific knowledge are key to solving complex problems in urban communities. In order to facilitate such an exchange, it is crucial to have fully-fledged privacy and security mechanisms in place.

The data that drives *Street Computing* poses some unique challenges in terms of privacy. It has a high potential for revealing personal information, "because urban sensing collects information in environments inhabited by and directly connected to human beings" (Cuff et al., 2008: 30). From mobile devices to cars or smart homes, users carry sensors on their bodies or have them installed on their property. As a result, sensing device custodians are engaging in constant self-surveillance. Furthermore, this monitoring may compromise the privacy of those who are not custodians or primary sensing targets – an effect that Campbell et al. liken to "second-hand smoke" (2008: 16). Finally, the collected data may reveal information that goes beyond the original purpose of the sensing system (Campbell et al., 2008; Cuff et al., 2008). Adams and Sasse aptly summarise this problem, stating that "most invasions of privacy are not intentional but due to designers inability to anticipate how this data could be used, by whom, and how this might affect users" (Adams & Sasse, 2001: 49).

Given the sensitivity of the information, it is unlikely that users will participate in *Street Computing* if they do not trust the underlying systems. One aspect of establishing trust is putting the users in charge of their personal data. This includes providing ways to control how personal data is collected, distributed, and processed. When sharing their data, users should have an option to become anonymous or pseudonymous and select the level of data granularity (Cuff et al., 2008). Designing

usable and effective user-interfaces for controlling privacy and security capabilities is challenging. One major reason is that "users value and want security and privacy functionality, but they regard this functionality as secondary to completing their primary tasks" (Karat, Karat & Brodie, 2005). Therefore, additional research in the area of Human-Computer Interaction is necessary to develop solutions that deliver a better user experience (Karat et al., 2005).

Different approaches have been proposed that allow people to make contributions to urban data collection efforts while providing strong individual privacy. One such solution is *PriSense*, which is based on data slicing and mixing and supports a wide range of statistical aggregation functions. Instead of directly answering a query, sensor nodes split their data into slices and pass them to randomly chosen cover nodes where they are mixed. From the responses of the different nodes, the initiator of the request can calculate an accurate aggregation result (Shi, Zhang, Liu & Zhang, 2010). A different approach is taken by *PoolView*, which is based on data perturbation. The idea is to make it impossible to reconstruct an individual data sequence's items or their trend without large error, while at the same time allowing to estimate the community data distribution and the average community data trend with high accuracy (Ganti, Pham, Tsai & Abdelzaher, 2008). The *MetroSense* project is addressing privacy issues by "providing sensing device custodians with a notion of anonymity through k-anonymous tasking" (Campbell et al., 2008; Cornelius et al., 2008).

4   Conclusion

One of the most profound changes to our world in recent times is the urbanisation of society. Since mid-2007, the majority of humankind lives in urban centres (Tibaijuka, 2008, p. 1-2), but this is only the beginning: a report from McKinsey & Company (2009) predicts that China alone will build 20,000 to 50,000 skyscrapers over the next twenty years; this is equivalent to ten current-day New York Cities. Australia is at the forefront of this global megatrend: it is forecast that by 2015, nearly 90% of Australia's population will dwell in urban areas (UNDP, 2008). These changes bring with them many challenges. Existing urban problems, including traffic congestion, pollution, stress on civic services, incidence of crime, etc. will intensify correspondingly (for the correlation between population size and crime volume and rate, e.g., see Nolan III, 2004). Yet the increasing population in cities affords numerous opportunities for research in the social sciences, architecture, urban planning and, of course, ICT. The Street Computing project sits squarely in this space and is uniquely positioned to produce significant results and tangible impact.

Rather than trying to develop this research agenda in isolation and with limited or non-disclosure contained within a single research lab, our aim right from the start has been to apply the "street" philosophy not only to the technology, but to our

development approach itself, too. For this reason we organised early on the aforementioned Street Computing workshop at OZCHI 2009 and shared our vision with likeminded colleagues working in the fields of urban informatics, human-computer interaction, and ubiquitous computing, and in turn encouraged them to share their own aspirations and related projects. We believe that this mutual exchange of works-in-progress will help to build momentum around an internationally significant and collaborative research agenda for urban technology and urban informatics. This special issue has been guest edited with this spirit in mind.

We are pleased that four teams of colleagues followed our invitation to revise their workshop submissions into full, peer-reviewed articles for this special issue of the Journal of Urban Technology. These contributions allow us to present and discuss broader aspects of the Street Computing agenda from multiple perspectives. **Vande Moere and Battino** focus on urban data visualisations. They introduce case studies that are based on both academic research projects as well as experimental design studios in order to postulate a set of design constraints for situated and public visualisation of urban data. **Flora Salim** argues that conventional ways to interrogate specific buildings or urban projects are often inadequate to deliver the micro-scale granularity required at this scale. She discusses the capacity of urban probes to conduct ambient and community sensing of environmental, behavioural or socio-cultural data that can inform urban design decisions. **Voss and Carolan** introduce the *Homesense* project which aims to bridge the gap between urban and domestic spaces as well as between purely technical and situated ("lived") expertise. The study initiated a participatory design approach for the user-led development of "open hardware toolkits" that were used as cultural probes. Finally, **Kloeckl, Senn and Ratti** of the Senseable City Lab at MIT present the *LIVE Singapore!* project. As far as we can tell, this project has the closest resemblance with the Street Computing vision, but with a focus on open data initiatives such as those that emerged from the Government 2.0 movement. *LIVE Singapore!* promises to become an open platform that developers interested in the City of Singapore can use to create innovative mash-ups of an increasing amount of urban data streams.

References

R. Ackoff, "From Data to Wisdom," *Journal of Applied Systems Analysis* 16:1 (1989) 3-9.

A. Adams, & M. Sasse, "Privacy in Multimedia Communications: Protecting Users, Not Just Data," paper presented at People and Computers XV: Interactions without Frontiers. Joint Proceedings of HCI 2001 and IHM 2001 (Lille, France, 2001, 49-64).

G. Bell, & P. Dourish, "Yesterday's Tomorrows: Notes on Ubiquitous Computing's Dominant Vision," *Personal and Ubiquitous Computing* 11:2 (2007) 133-143.

E. Ben-Joseph, H. Ishii, J. Underkoffler, B. Piper, & L. Yeung, "Urban Simulation and the Luminous Planning Table: Bridging the Gap between the Digital and

the Tangible," *Journal of Planning Education and Research* 21:2 (2001) 196-203.

S. Benford, M. Flintham, A. Drozd, R. Anastasi, D. Rowland, N. Tandavanitj, et al., "Uncle Roy All Around You: Implicating the City in a Location-Based Performance," paper presented at the ACM SIGCHI International Conference on Advances in Computer Entertainment Technology Singapore, 3-5 June 2004).

C. Bettini, O. Brdiczka, K. Henricksen, J. Indulska, D. Nicklas, A. Ranganathan, et al., "A Survey of Context Modelling and Reasoning Techniques," *Pervasive and Mobile Computing* 6:2 (2010) 161-180.

B. Biörnstad, & C. Pautasso, "Let It Flow: Building Mashups with Data Processing Pipelines," paper presented at the Service-Oriented Computing – ICSOC 2007 Workshops (Vienna, Austria, 17 Sep 2007, 15-28).

M. Boyle, & S. Greenberg, "The Language of Privacy: Learning from Video Media Space Analysis and Design," *ACM Transactions on Computer-Human Interaction* 12:2 (2005) 328-370.

O. Brdiczka, J. Crowley, & P. Reignier, "Learning Situation Models for Providing Context-Aware Services, " paper presented at the 4th International Conference on Universal Access in Human-Computer Interaction (Beijing, China, 22-27 July 2007, 23-32).

J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, et al., "Participatory Sensing," paper presented at the World-Sensor-Web Workshop, 4th ACM Conference on Embedded Networked Sensor Systems (ACM SenSys) (Boulder, Colorado, 31 Oct – 3 Nov 2006).

F. Calabrese, K. Kloeckl, & C. Ratti, "WikiCity: Real-time Location-sensitive Tools for the City," in M. Foth, ed., *Handbook of Research on Urban Informatics: The Practice and Promise of the Real-Time City* (Hershey, PA: IGI Global, 2009, 390-413).

A. Campbell, S. Eisenman, N. Lane, E. Miluzzo, R. Peterson, H. Lu, et al., "The Rise of People-Centric Sensing," *IEEE Internet Computing* 12:4 (2008) 12-21.

K. Cheverst, N. Davies, K. Mitchell, A. Friday, & C. Efstratiou, "Developing a Context-Aware Electronic Tourist Guide: Some Issues and Experiences," paper presented at Human Factors in Computing Systems (CHI 2000) (The Hague, The Netherlands, 1-6 April 2000, 17-24).

J. Chin, V. Callaghan, & G. Clarke, "An End User Tool for Customising Personal Spaces in Ubiquitous Computing Environments," in J. Ma, H. Jin, L. Yang & J. Tsai, eds., *Ubiquitous Intelligence and Computing. Lecture Notes in Computer Science Vol. 4159* (Berlin, Germany: Springer, 2006a, 1080-1089).

J. Chin, V. Callaghan, & G. Clarke, "An End-User Programming Paradigm for Pervasive Computing Applications," paper presented at the ACS/IEEE International Conference on Pervasive Services (Lyon, France, 26-29 June 2006, 325-328).

J. Corburn, *Street Science: Community Knowledge and Environmental Health Justice* (Cambridge: MIT Press, 2005).

C. Cornelius, A. Kapadia, D. Kotz, D. Peebles, M. Shin, & N. Triandopoulos, "Anonysense: Privacy-Aware People-Centric Sensing," paper presented at the

6th International Conference on Mobile Systems, Applications, and Services (Breckenridge, CO, 17-20 June 2008, 211-224).

CUD, *Connected and Sustainable Mobility / Connected Urban Development* (2009) <http://www.connectedurbandevelopment.org/connected_and_sustainable_mo bility/personal_travel_assistant> Accessed Jan 5, 2012.

D. Cuff, M. Hansen, & J. Kang, "Urban Sensing: Out of the Woods," *Communications of the ACM* 51:3 (2008) 24-33.

P. Dalsgaard, & K. Halskov, "Designing Urban Media Façades: Cases and Challenges," paper presented at Human Factors in Computing Systems (CHI '10) (Atlanta, Georgia, 10-15 April 2010, 2277-2286).

E. Della Valle, S. Ceri, D. Barbieri, D. Braga, & A. Campi, "A First Step Towards Stream Reasoning," paper presented at the Future Internet Symposium (Vienna, Austria, 28-30 Sep 2008, 72-81).

A. Dey, G. Abowd, & D. Salber, "A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications," *Human-Computer Interaction* 16:2 (2001) 97-166.

J. Erenkrantz, M. Gorlick, G. Suryanarayana, & R. Taylor, "From Representations to Computations: The Evolution of Web Architecture," paper presented at the 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (Dubrovnik, Croatia, 3-7 Sep 2007, 255-264).

G. Fitzmaurice, H. Ishii, & W. Buxton, "Bricks: Laying the Foundations for Graspable User Interfaces," paper presented at Human Factors in Computing Systems (CHI '95) (Denver, Colorado, 7-11 May 1995, 442-449).

M. Foth, R. Schroeter, & I. Anastasiu, "Fixing the City One Photo at a Time: Mobile Logging of Maintenance Requests," paper presented at OZCHI (Canberra, Australia, 28 Nov – 2 Dec 2011, 126-129).

R. Ganti, N. Pham, Y. Tsai, & T. Abdelzaher, "PoolView: Stream Privacy for Grassroots Participatory Sensing," paper presented at the 6th ACM Conference on Embedded Network Sensor Systems (ACM SenSys) (Raleigh, NC, 5-7 Nov 2008, 281-294).

R. Hague, P. Robinson, & A. Blackwell, "Towards Pervasive End-user Programming," paper presented at the Fifth International Conference on Ubiquitous Computing (Adjunct Proceedings) (Seattle, Washington, 12-15 Oct 2003, 169-170).

K. Henricksen, & J. Indulska, "Developing Context-Aware Pervasive Computing Applications: Models and Approach," *Pervasive and Mobile Computing* 2:1 (2006) 37-64.

K. Henricksen, & R. Robinson, "A Survey of Middleware for Sensor Networks: State-of-the-Art and Future Directions," paper presented at the International Workshop on Middleware for Sensor Networks (MidSens '06) (Melbourne, Australia, 28 Nov 2006).

J. Hightower, S. Consolvo, A. Lamarca, I. Smith, & J. Hughes, "Learning and Recognizing the Places We Go," paper presented at the 7th International Conference on Ubiquitous Computing (UbiComp '05) (Tokyo, Japan, 11-14 Sep 2005, 903).

S. Jordà, G. Geiger, M. Alonso, & M. Kaltenbrunner, "The reacTable: Exploring the Synergy between Live Music Performance and Tabletop Tangible Interfaces," paper presented at the 1st International Conference on Tangible and Embedded Interaction (TEI '07) (Baton Rouge, Louisiana, 15-17 Feb 2007, 139-146).

C. Karat, J. Karat, & C. Brodie, "Why HCI Research in Privacy and Security is Critical Now," *International Journal of Human-Computer Studies* 63:1-2 (2005) 1-4.

D. Kim, J. Hightower, R. Govindan, & D. Estrin, "Discovering Semantically Meaningful Places from Pervasive RF-Beacons," paper presented at the 11th International Conference on Ubiquitous Computing (UbiComp '09) (Orlando, Florida, 30 Sep – 3 Oct 2009, 21-30).

T. Kindberg, M. Chalmers, & E. Paulos, "Guest Editors' Introduction: Urban Computing," *IEEE Pervasive Computing* 6:3 (2007) 18-20.

V. Kulathumani, M. Sridharan, R. Ramnath, & A. Arora, "Weave: An Architecture for Tailoring Urban Sensing Applications across Multiple Sensor Fabrics," paper presented at the International Workshop on Mobile Devices and Urban Sensing (MODUS '08) (St. Louis, Missouri, 21 April 2008).

N. Lane, H. Lu, S. Eisenman, & A. Campbell, "Cooperative Techniques Supporting Sensor-Based People-Centric Inferencing," paper presented at the 6th International Conference on Pervasive Computing (Pervasive '08) (Sydney, Australia, 19-22 May 2008, 75-92).

D. Le-Phuoc, & M. Hauswirth, "Linked Open Data in Sensor Data Mashups," paper presented at the 2nd International Workshop on Semantic Sensor Networks (SSN '09), in conjunction with ISWC 2009 (Washington, DC, 26 Oct 2009).

M. Lee, Y. Uhm, Z. Hwang, Y. Kim, J. Jo, & S. Park, "A Ubiquitous Computing Network Framework for Assisting People in Urban Areas," paper presented at the 32nd IEEE Conference on Local Computer Networks (LCN '07) (Dublin, Ireland, 15-18 Oct 2007, 215-216).

E. Martinussen, J. Knutsen, & T. Arnall, "Bowl: Token-Based Media for Children," paper presented at Designing for User eXperiences (DUX '07) (Chicago, Illinois, 5-7 Nov 2007, 3-16).

McKinsey & Co., Preparing for China's Urban Billion (Beijing, China: McKinsey & Co., 2009) <http://www.mckinsey.com/Insights/MGI/Research/Urbanization/Preparing_for_urban_billion_in_China> Accessed Jan 5, 2012.

D. Merrill, J. Kalanithi, & P. Maes, "Siftables: Towards Sensor Network User Interfaces," paper presented at the 1st International Conference on Tangible and Embedded Interaction (TEI '07) (Baton Rouge, Louisiana, 15-17 Feb 2007, 75-78).

R. Murty, G. Mainland, I. Rose, A. Chowdhury, A. Gosain, J. Bers, et al., "CitySense: An Urban-Scale Wireless Sensor Network and Testbed," paper presented at the IEEE International Conference on Technologies for Homeland Security (Waltham, MA, 12-13 May 2008, 583-588).

M. Newman, A. Elliott, & T. Smith, "Providing an Integrated User Experience of Networked Media, Devices, and Services through End-User Composition,"

paper presented at the 6th International Conference on Pervasive Computing (Pervasive '08) (Sydney, Australia, 19-22 May 2008, 213-227).

J. Nolan III, "Establishing the Statistical Relationship between Population Size and UCR Crime Rate," *Journal of Criminal Justice* 32:6 (2004) 547-555.

J. Paay, J. Kjeldskov, S. Howard, & B. Dave, "Out on the Town: A Socio-Physical Approach to the Design of a Context-Aware Urban Guide," *ACM Transactions on Computer-Human Interaction* 16:2 (2009) 1-34.

L. Palen, & P. Dourish, "Unpacking 'privacy' for a networked world," paper presented at Human Factors in Computing Systems (CHI '03) (Fort Lauderdale, Florida, 5-10 April 2003, 519-523).

S. Patel, J. Rekimoto, & G. Abowd, "iCam: Precise at-a-distance Interaction in the Physical Environment," paper presented at the 4th International Conference on Pervasive Computing (Pervasive '06) (Dublin, Ireland, 7-10 May 2006, 272-287).

E. Paulos, R. Honicky, & B. Hooker, "Citizen Science: Enabling Participatory Urbanism," in M. Foth, ed., *Handbook of Research on Urban Informatics: The Practice and Promise of the Real-Time City* (Hershey, PA: IGI Global, 2009, 414-436).

E. Paulos, & T. Jenkins, "Urban Probes: Encountering Our Emerging Urban Atmospheres," paper presented at Human Factors in Computing Systems (CHI '05) (Portland, Oregon, 2-7 April 2005, 341-350).

P. Peltonen, E. Kurvinen, A. Salovaara, G. Jacucci, T. Ilmonen, J. Evans, et al., "It's Mine, Don't Touch! Interactions at a Large Multi-Touch Display in a City Centre," paper presented at Human Factors in Computing Systems (CHI '08) (Florence, Italy, 5-10 April 2008, 1285-1294).

P. Peltonen, A. Salovaara, G. Jacucci, T. Ilmonen, C. Ardito, P. Saarikko, et al., "Extending Large-Scale Event Participation with User-Created Mobile Media on a Public Display," paper presented at the 6th International Conference on Mobile and Ubiquitous Multimedia (MUM '07) (Oulu, Finland, 12-14 Dec 2007, 131-138).

B. Piper, C. Ratti, & H. Ishii, "Illuminating Clay: A 3-D Tangible Interface for Landscape Analysis," paper presented at Human Factors in Computing Systems (CHI '02) (Minneapolis, Minnesota, 20-25 April 2002, 355-362).

H. Raffle, C. Vaucelle, R. Wang, & H. Ishii, "Jabberstamp: Embedding Sound and Voice in Traditional Drawings," paper presented at the 6th International Conference on Interaction Design and Children (IDC '07) (Aalborg, Denmark, 6-8 June 2007, 137-144).

J. Rowley, "The Wisdom Hierarchy: Representations of the DIKW Hierarchy," *Journal of Information Science* 33:2 (2007) 163-180.

S. Seitinger, D. Perry, & W. Mitchell, "Urban Pixels: Painting the City with Light," paper presented at Human Factors in Computing Systems (CHI '09) (Boston, MA, 4-9 April 2009, 839-848).

J. Shi, R. Zhang, Y. Liu, & Y. Zhang, "Prisense: Privacy-Preserving Data Aggregation in People-Centric Urban Sensing Systems," paper presented at the 29th IEEE Conference on Computer Communications (INFOCOM '10) (San Diego, CA, 15-19 March 2010, 1-9)

J. Streefkerk, M. van Esch-Bussemakers, & M. Neerincx, "Field Evaluation of a Mobile Location-Based Notification System for Police Officers," paper presented at the 10th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '08) (Amsterdam, The Netherlands, 2-5 Sep 2008, 101-108).

A. Tibaijuka, Address at the 12th Session of the United Nations Conference on Trade and Development (UNCTAD XII) by Anna K. Tibaijuka, UN Under-Secretary-General and Executive Director, UN-HABITAT and Director-General, UN Office at Nairobi (UNON) (Accra, Ghana: UN-HABITAT, 20-25 April 2008) <http://www.unhabitat.org/content.asp?cid=5684&catid=649&typeid=8> Accessed Jan 5, 2012.

K. Truong, E. Huang, & G. Abowd, "CAMP: A Magnetic Poetry Interface for End-User Programming of Capture Applications for the Home," paper presented at the 6th International Conference on Ubiquitous Computing (UbiComp '04) (Nottingham, UK, 7-10 Sep 2004, 143-160).

B. Ullmer, & H. Ishii, "mediaBlocks: Tangible Interfaces for Online Media," paper presented at Human Factors in Computing Systems (CHI '99 extended abstracts) (Pittsburgh, Pennsylvania, 15-20 May 1999, 31-32).

B. Ullmer, & H. Ishii, "Emerging Frameworks for Tangible User Interfaces," *IBM Systems Journal* 39:3-4 (2000) 915-931.

UNDP, Human Development Report 2007 / 2008 (New York, NY: United Nations Development Programme, 2007) <http://hdr.undp.org/en/media/HDR_20072008_EN_Complete.pdf> Accessed Jan 5, 2012.

J. Underkoffler, & H. Ishii, "Urp: A Luminous-Tangible Workbench for Urban Planning and Design," paper presented at Human Factors in Computing Systems (CHI '99) (Pittsburgh, Pennsylvania, 15-20 May 1999, 386-393).

J. Underwood, H. Smith, R. Luckin, & G. Fitzpatrick, "E-Science in the Classroom – Towards Viability," *Computers & Education* 50:2 (2008) 535-546.

M. Weiser, "The Computer for the 21st Century," *Scientific American* 265:3 (1991) 94-104.

K. Whitehouse, F. Zhao, & J. Liu, "Semantic Streams: A Framework for Composable Semantic Interpretation of Sensor Data," paper presented at the Third European Workshop on Wireless Sensor Networks (EWSN '06) (Zurich, Switzerland, 13-15 Feb 2006, 5-20).

R. Wishart, K. Henricksen, & J. Indulska, "Context Obfuscation for Privacy via Ontological Descriptions," paper presented at the First International Workshop on Location- and Context-Awareness (LoCA '05) (Oberpfaffenhofen, Germany, 12-13 May 2005, 339-353).

J. Wong, & J. Hong, "Making Mashups with Marmite: Towards End-User Programming for the Web," paper presented at the Human Factors in Computing Systems (CHI '07) (San Jose, CA, 28 April – 3 May 2007, 1435-1444).

F. Zambonelli, "Pervasive Urban Crowdsourcing: Visions and Challenges," paper presented at the IEEE International Conference on Pervasive Computing and Communications (PerCom '11) (Seattle, WA, 21-25 March 2011, 578-583).

C. Zhou, D. Frankowski, P. Ludford, S. Shekhar, & L. Terveen, "Discovering Personally Meaningful Places: An Interactive Clustering Approach," *ACM Transactions on Information Systems* 25:3 (2007) Article 12, 1-31.

O. Zuckerman, T. Grotzer, & K. Leahy, "Flow Blocks as a Conceptual Bridge between Understanding the Structure and Behavior of a Complex Causal System," paper presented at the 7th International Conference on Learning Sciences (ICLS '06) (Bloomington IN, 27 June – 1 July 2006, 880-886).