Smart Services Across the Real and Virtual Worlds

Eleni Stroulia
stroulia@cs.ualberta.ca
Department of Computing Science
University of Alberta

Abstract

Today, we are witnessing a new level of scale, complexity, and pervasiveness of software systems that are designed to support much more holistically complex processes. Much richer information is becoming available for processing, including raw data, proven facts and people's opinions. A multitude of interesting analytics methods are being developed to extract new knowledge from this plethora. In this paper, we examine some opportunities for smart service delivery that arise from the integration of sensor networks, service-oriented middleware and virtual worlds.

1 "Smart Systems"

Today, we are witnessing a new level of scale, complexity, and pervasiveness of software systems. Tiny sensors, embedded in our environment, houses, clothes, and bodies, measure an unprecedented variety of environmental and physiological parameters. Social-networking platforms enable practically everyone to socialize online and publish their knowledge and opinions, giving rise to a multitude of flexible communities, through which information is fast propagated. And, through high-level programming languages and cross-platform integration standards, legacy and new software systems are being integrated, increasing the efficiencies of complex processes (that can now be globally optimized) and generating opportunities for new services.

This constellation of technologies have brought about the vision of a new breed of systems that will support *smart services* in a whole slew of domains, including education, health care, utilities, retail and manufacturing. Although the term "smart system" is by no means well defined, it is usually (see http://asmarterplanet.com/) associated with three attributes: (a) instrumented, (b) intelligent and (c) interconnected.

The term "***instrumented***" refers to the embedding of "sensing devices" in the environment where the service-delivery process occurs. Sensing devices vary widely in terms of their costs and capabilities. RFIDs (passive and active) are being increasingly embedded in objects for warehouse management and logistics applications. Sensors deployed through wireless networks are envisioned as the smart dust covering our planet and monitoring a multitude of environmental parameters. At the same time, wireless transceivers are being embedded in all types of electronic devices, making available their remote monitoring and control. Many cell phones today include accelerometers, GPS and compass, enabling a rather precise recognition of where someone is located, whether they are moving, how fast, and in what direction.

This large-scale instrumentation of the world around us is also contributing to the feature of ***interconnectedness***. On one hand, multiple sensing devices collect different measurements of the same phenomena: for example, a person walking through a mall may be sensed through his cellphone GPS as well as the Bluetooth-enabled store devices that recognize him as a loyal customer to email him offers. On the other, measurements collected for different purposes can potentially be analyzed together: for example, drug dispensation usage (collected primarily for insurance reimbursement purposes) can be analyzed together with at-home drug usage (monitored by smart pillboxes) and physiological parameter recordings (collected by home care devices) to produce information about patients' conformance with prescribed drug usage and the drugs' clinical

---

effectiveness. Another aspect of interconnectedness is the large scale of intra- and inter-organization integration of hardware and software. Through hardware virtualization and task distribution heterogeneous computing resources appear as a single computing platform; and through web-service technologies legacy and new software is being integrated to support old and new processes. The last aspect of interconnectedness refers to the social networks supported by the variety of social applications available today. Each individual may belong to a multitude of online social networks as a result of their membership to real-world organizations or online communities of interest or the more entertainment-oriented virtual worlds. These densely populated, highly overlapping networks are highly effective channels for information dissemination, annotation (with subjective opinions) integration and cross-pollination.

Finally, the term "***intelligent***" refers to the numerous analyses – both computational and interactive – that are possible on the multitude of information available. Analytical and interactive data mining can enable the fusion of different types of data towards higher-order information that can support better – more informed, more context specific, more timely – decision making.

In our own recent work, we have been developing software to support several "traditional" service areas (including health care, education, and more recently manufacturing). In this work, smart systems are envisioned to support the whole service-delivery process. To that end, (a) they are built around a core service-oriented middleware that integrates systems (semi)automating individual process steps; (b) they include a wireless sensor-network that senses their environment and/or mobile devices through which users manipulate information as they perform their tasks in the context of the overall process; and (c) they incorporate social platforms (2D wikis and 3D virtual worlds) to support the natural interactions of the people involved in the service-delivery process.

2 Service-Delivery across the Real and Virtual Worlds

To date, we have worked on three distinct smart systems that exemplify three types of integration between the real and virtual worlds, which we review below.

2.1 The Smart Condo

In this project, the virtual world serves as a mirror of the real world.

The service: The Smart Condo (Boers et. al 2009; Stroulia et al. 2009) is a collaborative project with colleagues in Networks, Rehabilitation Medicine, Art and Design, Pharmacy, and Education. The clinical objective is to support older adults and patients living at home or in rehabilitation facilities. Although these individuals are, by and large, able to live independently, they are still susceptible to harmful incidents related to physical infirmities and/or cognitive impairments. To that end we are designing a model Smart Condo in which patients can live assisted by software and health-care providers who can offer their services in an unobtrusive yet timely manner.

The system: The physical space is designed according to Universal Design principles (http://www.design.ncsu.edu/cud/about_ud/udprinciples.htm). Sensors have been embedded in many of its architectural elements (motion sensors on walls, accelerometers on doors), appliances (switches on microwave oven) and furniture (pressure sensors on chairs and beds and switches on cabinet doors) and a wireless network collects their readings in a server. We are now working towards integrating specialized home-care devices, such as a smart pillbox (alerting the patient when it is time to get a medicine and recording when the patient does so) and a watch that monitors blood pressure and heart rate. The incoming stream of raw readings is mined to recognize patterns indicative of interesting activities in the condo; for example, combination of motion sensors are triangulated to recognize a person's location in the condo. Further analyses are applied on the collected data; for example, through a sequence of recorded locations a path planner is invoked to infer a complete trajectory of how the person has moved in the space.

The server exposes this information through a set of REST (Fielding and Taylor 2002) APIs to a variety of possible clients. We have developed (a) interactive plots of sensor readings through which users can explore data trends and statistics, and (b) a wiki through which they can annotate interesting reading collections. The most interesting client is a reengineered Second Life client that accesses information regarding the patient's activities and uses it to control an avatar within a model of the Smart Condo SL, mirroring the patient's activity in the real-world Smart Condo.

How it is smart: The system integrates a few tools that make its deployment straightforward, including support for (semi)automatically building the home's 3D model from its architectural drawings, and an interactive tool for deciding the sensors' placement. In this manner, we were able to deploy it within several hours in a new space.

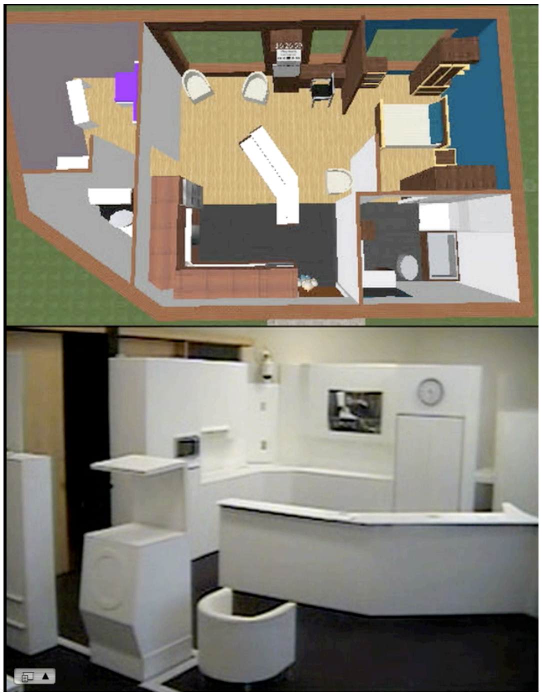

Figure 1: Real and Virtual Smart Condo

When compared with video, its natural "competitor", the virtual-world view into the patient's activity is realistic and intuitive, more cost-effective in terms of required network bandwidth, while at the same time less intrusive, since personal-appearance details are not actually monitored or recorded. In this manner, healthcare professionals have a live stream of the person's activity and can be alerted to intervene at the occurrence of a harmful event. They can even join the patient in their virtual home to answer questions and offer advice, thus associating an added value to the cost of the intrusion to the patient's privacy.

More importantly, the Smart Condo consistently maintains the patient's history, relevant to many clinical decisions, enabling them to improve their practice by infusing it with more evidence. For example, our colleagues from pharmacy are interested in correlating drug-dispensation data (available at the pharmacy) with drug usage and physiological-parameter values, in order to assess the patient's compliance with their drug regimen and the effectiveness of this regimen. On the other hand, our rehabilitation-medicine colleagues are interested in examining how the environment affects the level and type of activity of their patients, in order to advance proper design standards and building codes and to come up with specialized rehabilitation regimens for their patients appropriate for their environments.

2.2 Medical Simulation Training

In this activity, the virtual world is viewed as an (equivalent) alternative to the real world.

The service: Our work on medical-simulation training in a virtual world is in collaboration with colleagues from Education, Nursing and Emergency Medical Services. The stabilization of a trauma patient by EMS personnel at the scene of the accident and the handoff of the patient to the ER personnel is a complex procedure. It has to conform with many rules and meet timing constraints and it involves the handling of a variety of devices and the coordinated interactions among several people. Training students – across health disciplines – for this complex process is very expensive; it is usually done with standardized-patient scenarios in simulated scenes where the role of the patient is enacted by a professional actor. A virtual-world based simulation can potentially enable the training of many more students by providing a lower-cost alternative while sacrificing little of the verisimilitude of the experience.

The system: The MixEd Reality Integrated Training System (MERITS) framework (Chodos et al. 2009) integrates a web-services enactment engine with a virtual world. Together, these two components model the complete EMS/ER scenario. The prescribed procedure, including the victim assessment and stabilization activities that take place at the scene and the information communication activities among the people involved, is documented in a set of structured Annoki pages (Bauer et al. 2009). It is then "translated" by software analysts into a BPEL model. The student participants, the artifacts involved in the scenario (including the accident and ER scenes and the medical devices) and the specific actions that the participants may carry out (including manipulating the victim, using the medical devices, and coordinating with each other) are modeled with corresponding avatars, virtual objects and their behavioral scripts. These scripts are all equipped with a SOAP module that communicates to the BPEL-enactment engine information of interest about the state of the virtual world and the participants' behaviors. In this manner, the BPEL engine can recognize the progress of the student participants and potentially issue warnings to them, if their behaviors do not conform to the specified procedure.

How it is smart: The system provides support to an important problem in developing training scenarios, namely their specification. Annoki – our suite of MediaWiki extensions – provides a collaborative repository for easy prototyping of the scenario scenes. The fundamental contribution of this work will eventually be its broad enablement of simulation-based training. Today, simulations presuppose the synchronous availability of a skilled professional, as the "director", and a set of costly resources, including the venue, all the required props and all the required participants. This fundamentally limits the adoption of simulations to "rich" organizations, few typical scenarios, and sometimes within a single discipline (to limit the number of participants). High-quality virtual-world simulations with tools to ease the scenario specification for domain experts will enable the treatment of many (including rare) scenarios, the participation of multiple, geographically distributed stakeholders from across disciplines.

Figure 2: Accident Scene in the EMS/ER Scenario

The current application of MERITS is the simulation-based training of health professionals. In the longer run, however, our objective is to investigate the novel breed of simulation-based training that it enables: namely, training for complex *mixed* service-delivery processes, where some activities are carried out by people and some are automated. Note that these type of processes abound in all service-delivery domains and it is a long-standing objective of service scientists – from business, computing and operations research alike – to further systematize them and increase the degree of their automation (Zysman 2006). The MERITS orchestration of automated web-services with people's activities in the (simulation of the) real world enables a new style of service-process modeling and auditing. Interactive and computational analyses can be applied to the process-instance traces to identify opportunities of their improved management and potential reengineering.

2.3 Education AARGs

In our nascent augmented/alternative reality education games (AARG) work, the virtual world *augments* the real world. Different but correlated (inter)actions are possible in the real and the virtual world with consequences across both worlds.

The service: In our recent collaboration with colleagues from Humanities Computing, we are developing a technological infrastructure for creating campus-wide collaborative, educational, augmented-alternate reality games (AARGs). The game players will interact with the game (a) in the real world through location-specific clues (communicated to them through cell-phones, PDAs, and later sensors embedded in the environment) and (b) in a parallel virtual world that reflects the real world in some dimensions (currently space) and augments it in others (currently providing the design tools for designing and placing artificial characters in the game).

The system: The core middleware in this project is a web portal, likely implemented on top of Annoki, our own extension to MediaWiki. The portal services are similar to that of a social-network application, where players publish in their profiles information about themselves and their game statistics, and communicate with other players to discuss their successes and exchange game clues. The virtual world will likely be Second Life or simply a GIS of the campus map and will be used primarily to place virtual characters in the real world, as virtual "geo-cached" objects.

How it is smart: We envision two distinct styles of AARG game narratives. In one style, virtual-world play will allow participants to interact digitally with the environment, allowing the impact of gaming decisions to be simulated (in super-real time) in the virtual world in order to evaluate their consequences. In this scenario, it is the power of the analytical models simulating the impact of the players' actions and the realism of the virtual world effectively visualizing the simulation results that makes the system smart.

In the latter style, which we explore in the current game, real-world play will be augmented through displays synchronized with the virtual world. Tablets or PDAs with realistic overlay technologies (i.e., http://layar.com/ and http://cenceme.org/) and projection eye-ware will overlay the virtual world with the real world. In this scenario it is the creativity of the overlay and the degree to which it engages the players' imagination that makes the system compelling.

3 Research Challenges

In these three activities, we have faced several problems; some of them technical, most of them conceptual. In this section we review, very briefly, what we consider the most intellectually interesting ones.

Service modeling: Acknowledging and modeling the complexity of service processes is essential to supporting them with smart systems. To date, there is no conceptual service-modeling framework that accounts for all aspects of social interactions, technology and physical environment parameters around a service process. In our work with MERITS, we have started examining the service-blueprinting methodology (Bitner et al. 2007). Service blueprinting concerns itself with (a) customer actions, (b) onstage/visible contact employee actions, (c) backstage/invisible contact employee actions, (d) support processes, and (e) physical evidence. This framework is broad enough but very informal. To increase its utility, it should be integrated with (i) more traditional software models of workflows (to express the orchestration of the various onstage/backstage activities), (ii) artificial-intelligence models of behavior (to represent the behaviors of the participants' and their avatars in the virtual world), and (iii) design models of physical space (to specify the layout and affordances of the space in which the service is delivered and to represent it in the virtual world).

Separating the layers: The integration of traditional middleware with wireless sensor networks, mobile devices and highly interactive environments such as virtual worlds poses substantial software-engineering challenges. The elementary functionalities that each layer delivers and the orchestration of these functionalities are not yet well understood and delineated. For example, a wireless sensor network can be viewed either as a (set of) untrustworthy device(s) collecting raw data, or as a computational component that collects data while also computing interesting higher-order queries on them. Similarly, the behaviors of avatars in a virtual world have information processing and communication consequences of interest to the traditional workflow middleware. The question now becomes to precisely characterize the layers within which these functionalities can be organized and the APIs through which they can be composed.

Managing identities: In such complex multi-component systems, the user's identity is a function of her web profiles and virtual-world avatars, tool-usage history and in-world activities, and real-world behavior (as captured by wearable sensors, including mobile devices). In order to make possible the consideration of all these sources of user information within the same context, we have to develop a transparent and flexible mechanism for managing the user's identity across platforms and systems. In our work, we have been associating the various IDs of a single user with her MediWiki's "user" specification. We believe (Bauer et al. 2009) that wikis offer an easy to grasp metaphor for large-scale collaboration and social networking and that MediaWiki is a flexible platform for integrating resources and supporting access to them through REST APIs. On the other hand, the presentity specification (http://tools.ietf.org/html/rfc2778) is aimed at capturing a user's presence at the various channels through which he may be accessible.

Information fusion: More platforms are becoming available for information representation. The real world "generates" information through sensor readings. The web contains factual information and opinions accessible through services. In 3D virtual worlds, aspects of the real world are modeled and simulated. The challenge now becomes to recognize the semantic overlap among these information sources, to use its redundancy to increase the robustness of all tasks that rely on it, and to fuse the available multi-modal data (i.e., data of different types) on the same process into meaningful evidence for understanding it better. For example, in the Smart Condo project, we can imagine correlating information about the sensed patient's activity with his blog (assuming there is one) and with the comments/finding of their health providers, in order to extract high-quality evidence on their status and better manage their care.

Knowledge Mining: Much of the intelligence of these systems relies on analytics. Today's distinction between data mining (where knowledge is extracted through computationally intensive processes on a database), OLAP (where users interactively manipulate and search through data sets) and stream mining (where data is analyzed in the context of other data available in the same temporal window) is rather crude. Finer grained categories are required to appropriately characterize the quality of the available input data (in terms of precision and recency), the quality of the output data required to appropriately support a specific activity or decision, and the cost of the computations or user interactions that can potentially accomplish this transformation.

Explicit workflows vs. emergent interactions: Some service-delivery processes are more explicitly orchestrated than others; consider, for example, the activities of a cashier vs. a personal sales advisor. When there is no explicit "workflow" to drive the participants' interactions, it is the environment that has to engage them and guide them with appropriate cues to make progress towards their goals. To that end, we need to better understand the cognitive mechanisms of motivation and information filtering in people, and to develop methods for endowing our systems with the capacity to motivate their users.

4 Concluding Remarks

Imagining the future of our computational systems and their embedding in our environment and life is ultimately exciting. In this paper, we have discussed the three styles of real-and-virtual service systems and the services they support (healthcare delivery, professional procedural training, and learning). We have also briefly mentioned some of the most interesting challenges we have faced in this work until now. Undoubtedly, these system styles will soon be viewed as just a few points in a complex space of smart systems. And the challenges put forward will be refined and redefined as the work of the research community progresses. Nevertheless, we believe that the constellation of these technologies (sensor, service-oriented middleware and virtual worlds) has huge potential to support smart services and is rich with many interesting problems to the research community.

Acknowledgements

I would like to acknowledge the generous financial support of iCORE, NSERC, and IBM. More importantly, I am grateful to my students (N. Boers, D. Chodos, V. Ganev, K. Golmohammadi, J. Huang, P. Kuras, P. Naeimi, D. Serrano, M. Smit) and my colleagues (Drs. Boechler, Carbonaro, Gburzynski, Gouglas, Guirguis, King, Lederer, Liu, Messinger, Nikolaidis, Rockwell, Sadowski, Schaeffer, and Mrs. Tansey, Reid and Brisbin) for their intellectual curiosity, enthusiasm, sense of humor, and hard work.

About the Author

Eleni Stroulia is a Professor and iCORE Industrial Research Chair with the Department of Computing Science at the University of Alberta. Her research addresses industrially relevant software-engineering problems with automated methods, based on artificial-intelligence techniques. Her team has produced automated methods for migrating legacy interfaces to web-based front ends, and for analyzing and supporting the design evolution of object-oriented software. More recently, she has been working on the development, composition, run-time monitoring and adaptation of service-oriented applications, and on examining the role of Web 2.0 tools and virtual worlds in enabling the practices of collaborating communities.

References

[1] K. Bauer, M. Fokaefs, B. Tansey and E. Stroulia: WikiDev 2.0: Discovering Clusters of Related Team Artifacts. CASCON 2009, Toronto, Canada, November 2-6, 2009.

[2] M. J. Bitner, A. L. Ostrom, F. N. Morgan: Service Blueprinting: A Practical Tool for Service Innovation, 2007.

[3] N. Boers, D. Chodos, J. Huang, P. Gburzynski, I. Nikolaidis, E. Stroulia: The Smart Condo: Visualizing Independent Living Environments in a Virtual World, Pervasive Health 2009.

[4] D. Chodos, P. Naeimi, E. Stroulia: A simulation-based training framework for health-science education on video and in a virtual world, 2(1) Journal of Virtual Worlds Research. 2009

[5] R. T. Fielding, R. N. Taylor: Principled design of the modern Web architecture. ACM Trans. Internet Techn. 2(2): 115-150 (2002)

[6] E. Stroulia, D. Chodos, N. Boers, J. Huang, P. Gburzynski, I. Nikolaidis: Software Engineering for Health Education and Care Delivery Systems: The Smart Condo Project, workshop on "Software Engineering in Healthcare", at the 31st International Conference on Software Engineering, May 2008, Vancouver Canada, pp. 20-28, ACM/IEEE.

[7] J. Zysman: The algorithmic revolution – the fourth service transformation. Communications of ACM 49(7): 48 (2006).