Proceedings

First Symposium on the Personal Web

November 3, 2010

A Message of Welcome

The First Symposium on the Personal Web Chairs

It is our great pleasure to welcome each of you to this, the First Symposium on the Personal Web, sponsored by the IBM Canada Centre for Advanced Studies Research and co-located with CASCON 2010 in Markham, Ontario. The purpose of this symposium is to bring together prominent researchers and practitioners from a diverse range of research areas whose expertise can be leveraged in the advancement of science and practice relating to the Personal Web. Research on the Personal Web is an outgrowth of the Smart Internet initiative, which seeks to extend and transform the web to be centered on the user, with the web as a 'calm' platform ubiquitously providing cognitive support to its user and their tasks. As with the preceding SITCON workshop (held at CASCON 2009), this symposium involves a multi-disciplinary effort that brings together researchers and practitioners in data integration; web services modeling and architecture; human-computer interaction; predictive analytics; cloud infrastructure; semantics and ontology; and industrial domains such as health care and finance.

The goal of the symposium is to discuss and investigate different aspects of the architecture and functionality needed to make the Personal Web a reality. Interactions during the symposium should identify opportunities for linking research and sharing results. The longer term outcome includes the creation of an edited book with the tentative title "The Personal Web" based on extended versions of the symposium proceedings. Another anticipated outcome is the deployment of a public instance of the Personal Web infrastructure in the IBM cloud in order to generate real life data that can be used in future studies. The goal of that "venture research" will be to collect data that can guide future directions in the Personal Web, with researchers amending data models, ontologies and functionality as appropriate to match real-world practice. The symposium will be held over the course of a full day, being launched by a brief welcome introduction, followed by a keynote presentation by Joanna Ng, introducing the Personal Web vision and related research challenges. A series of twenty papers will then be presented covering a wide range of research areas relating to the Personal Web, followed by a closing session of discussion. Each presenter will be encouraged to present one major research question or challenge relating his or her area of interest to the Personal Web.

We would like to thank those who submitted papers to these proceedings, along with all the symposium participants, for bringing their expertise and enthusiasm to this event. This symposium would not have been possible without the generous support of the IBM Canada Centre for Advanced Studies Research, and we are grateful for their support. A number of people have put a lot of effort into organizing this workshop, and we would like particularly to thank Emilia Tung, Debbie Kilbride, and Jimmy Lo for their help. We look forward to active and enjoyable discussion on this important and exciting topic.

The Symposium Chairs:

Mark Chignell, Department of Industrial Engineering, University of Toronto
James R. Cordy, School of Computing, Queen's University
Joanna Ng, IBM Canada Centre for Advanced Studies Research
Yelena Yesha, University of Maryland, Baltimore County

Proceedings of the First Symposium on the Personal Web

Co-located with CASCON - Markham, Ontario, Canada, Nov.3, 2010
Sponsored by the IBM Canada Centre for Advanced Studies Research

Table of Contents

1. Symposium Introduction

First Symposium on the Personal Web
Mark Chignell University of Toronto, James R. Cordy, Queen's University, Joanna Ng, Center for Advanced Studies, IBM Canada, Yelena Yesha, University of Maryland, Baltimore County

2. Keynotes and Visions of the Personal Web

Keynote: The Personal Web
Joanna Ng, IBM Canada Laboratory, Toronto, Ontario, Canada

Modeling and Analysis of Personal Web Applications: A Vision
Marsha Chechik, Jocelyn Simmonds, Shoham Ben-David, Shiva Nejati, Mehrdad Sabetzadeh, and Rick Salay (Shiva and Mehrdad are with Simula Research Lab, Norway. Other authors are with Department of Computer Science, University of Toronto)

The Personal Web: Personal Observations and the Challenge for Financial Services
Ian Graham, Bank of Montreal

3. Managing Context

Notification Design Challenges for the Personal Web
Jin Li and Jimmy Lo, IBM Canada

Managing Dynamic Context to Enable User-Driven Web Integration in the Personal Web
Norha M. Villegas, Hausi A. Müller University of Victoria; Joanna Ng, Alex Lau IBM Canada

Context-Aware Personal Web Portals
Jens Weber-Jahnke, University of Victoria

4. Analytics and Visualization

User, Social and Location Context for IBM Business Analytics
Stephan Jou, IBM Canada

Towards Smart Visual Data Exploration Tools
Lars Grammel and Margaret-Anne Storey, University of Victoria

Supplementing Semantics with Statistics in the Personal Web
Mark Chignell and Ryan Kealey, University of Toronto

5. Healthcare

Information Feudalism: How Legacy Information Systems Kill People in the Healthcare System
Jacques Lee, Sunnybrook Health Sciences Centre

The Personal Web and Clinical Medicine: A Physician's View
Tammy Sieminowski, Bridgepoint Hospital

Smart Interactions for Health Self-Management
Norm Archer, McMaster University

Smart Health and Wellbeing
Yelena Yesha, University of Maryland Baltimore County

Toward Patient-centric Healthcare
Igor Jurisica, Ontario Cancer Institute

6. Sociality and Privacy

Persons, Communities and Artifacts on the Web
Eleni Stroulia, University of Alberta

Paying Attention in Meetings: Multitasking in Virtual Worlds
Kelly Lyons, University of Toronto

Towards Smart Privacy on the Personal Web
Reza Samavi and Mariano Consens, University of Toronto

6. Personalized Services

A Framework for Composing Web Resources
Hua Xiao, Bipin Upadhyaya, Ran Tang, Ying Zou, Queen’s University; Joanna Ng and Alex Lau, IBM Canada

Service Subscription for Personal Web
Chunyang Ye and Hans-Arno Jacobsen, University of Toronto

Web Personalizer as User Consultant
Kamran Sartipi, McMaster University

WSCells for the Personal Web
Douglas Martin, James R. Cordy, Queen’s University

Symposium Introduction

First Symposium on the Personal Web

Mark Chignell, Univ. of Toronto, Toronto, Ontario
James R. Cordy, Queen's University, Kingston, Ontario
Joanna W. Ng, IBM Canada Laboratory, Toronto, Ontario
Yelena Yesha, Univ. of Maryland, Baltimore County

Abstract

The First Symposium on the Personal Web is co-located with CASCON 2010 and is sponsored by IBM CAS Research. This symposium includes key researchers and practitioners in a range of related areas in order to organize and focus the research directions and challenges of the Personal Web as the next instantiation of the Smart Internet. The smart internet is envisaged as a platform for automatic, dynamic aggregation of data and services for the purpose of supporting each user's goals; tasks and concerns, both cognitively and socially. The Personal Web focuses on the user view of the Smart Internet, allowing users to conduct ad hoc or persisted integration effortlessly across the web, as per their context and spheres of interest.

The Smart Internet research initiative has two distinct areas of research. *Smart Interactions* address factors that impact the discovery; aggregation and delivery of data and services from the internet that are most relevant and appropriate to the users' situations and tasks at hand. *Smart Services* address the challenges in the underlying web architecture and runtime infrastructure as the behind-the-scene enabler to actually deliver the data and services in a manner that meet the requirement for the support of smart interactions. Building on what has been previously established in Smart Internet, the Personal Web emphasizes the use of *semantically-linked data* as its fundamental basic building block in order to enable user-sovereign, open integrations across the web, and to achieve the goals of users with minimal cognitive effort. This refines the research scope of smart interactions and smart services, and creates new problem statements and research challenges.

1 Rationale

In the CASCON 2009 workshop on Smart Internet, which we refer to as SITCON, we focused on developing and refining the concepts and requirements of the Smart Internet, including Smart Interactions and Smart Services. We also proposed technologies and methodologies in enabling computing needed for real world implementations.

Since the SITCON workshop at CASCON 2009, subsequent refined research directions, innovations and research results have been captured in an edited book on the Smart Internet to be published by Springer in fall 2010 [1].

The evolving goals of this research initiative have resulted in this first Symposium of the Personal Web, positioning the notion of Personal Web as an instantiation of the Smarter Internet. The plan is to focus on bridging the current technology gaps in the realization of the goal of empowering "ME", as a general user, to control and conduct *open*, *user sovereign web aggregation* for ad hoc or persisted integration across the web. The Personal Web should be implemented with such simplicity in user experience that even non-technical users can have a rich experience of personalized integration through various kinds of integrate-able web elements (such as web domains' structural data and restful web services), data and restful web services), made possible by web software programmers.

As with the preceding SITCON workshop, this symposium will involve a multi-disciplinary effort that brings together researchers in data integration; web services modelling and architecture; human computing interaction; predictive analytics; cloud infrastructure; semantics and ontology; as well as industrial domains such as health care; journalism; e-Commerce and others.

The purpose of the symposium is to bring together researchers and practitioners from these multi-disciplinary areas in order to organize and focus research directions and challenges and to brainstorm solutions for the advancement of related science needed to enable real life deployment of the Personal Web.

2 The Personal Web

This Symposium on the Personal Web is a continuation of the Smart Internet research initiative. The concept of the Personal Web was originated and proposed by symposium co-chair Joanna Ng, with relevant research by IBM Canada CAS Research affiliated university faculty members being discussed in this first symposium.

The Personal Web gives the general user the power and control to assemble and aggregate what the user is interested in across the web; including contextually relevant and personal data, as well as relevant web services as modules of capability, without the need for assistance from a programmer. In the Personal Web, multiple web domains are melded into integrated views and interactions involving visualization and web artifacts, all driven by the user's context and spheres of interest; making the web platform appear unified. Personal Web interaction will serve the user's situational and/or persistent needs for integration in a manner that is cognitively supportive, and simple enough for everyday use by non-technical users.

The Personal Web is currently envisioned in terms of three main areas. These are the areas of Smart Interactions and Smart Services the comprise the Smart Internet, but with the area of Smart Services further broken down into Smart Services for Data and Smart Services for Tasks. Specific research areas and tasks that will be relevant to the research activities discussed in the workshop include:

- Enabling existing web domains for participation in the Personal Web
- User Data and Knowledge Management for Personal Web Integration
- Elicitation of Personal Web Knowledge (e.g., through Analytics or Inference)
- Data and Service Retrieval
- Data and Services Integration and Classification
- Data and Services Interaction and Visualization
- Personal Web User Interfaces and User Experiences
- Privacy and security for Personal Web
- Semantics and Ontology for Personal Web services and data
- Semantic relationships and links discovery

3 Organizers

The co-chairs for the First Symposium on the Personal Web will be Mark Chignell, Professor of Mechanical and Industrial Engineering at the University of Toronto, James Cordy, Professor at Queen's University's School of Computing, Joanna Ng, Head of IBM Canada's Centre for Advanced Studies Research and Yelena Yesha, Professor of Computer Science and Electrical Engineering at the University of Maryland, Baltimore County.

4 Format

The Symposium will have three distinct sessions. In the first session, invited participants will be presented with an overview of the research initiative of the Personal Web to lay out the scope, including the key research challenges and problem statements and why collaboration is necessary and how the collaboration could take place.

The second session involves a set of presentations to share research ideas and different research challenges on three major sub-topics, namely *Smart Interactions, Smart Services for tasks and Smart Services for data.* These three groups of presentations arise from IBM's experience with challenges in these areas and are subject to revision based on input from both academic participants and companies collaborating in this initiative.

The third session is a panel discussion on the Personal Web. The panel discussion will cover the mission and goals of the Personal Web along with the key research areas needed to enable it.

5 Expected Outcomes

The goal of the symposium is to bring together Canadian researchers who are investigating different aspects of the architecture and functionality needed to make the Personal Web a reality. Interactions during the symposium should identify opportunities for linking research and for sharing results. The longer term outcome includes the creation of an edited book out of a revised version of the symposium proceedings with the tentative title "The Personal Web". Another anticipated outcome is deployment of a public instance of the Personal Web infrastructure in the IBM cloud in order to generate real life data that can be used in future studies. The goal of that "venture research" will then be to create realistic data that can guide future directions in the Personal Web, with researchers amending data models, ontologies, and functionality as appropriate to match real-world practice.

Reference

[1] Chignell, M., Cordy, J., Ng, J. and Yesha, Y. (Eds). *The Smart Internet*. Lecture Notes in Computer Science, 6400. Berlin: Springer-Verlag.

Keynotes

&

Visions

of the Personal Web

The Personal Web
Smart Internet for Me

Joanna Ng
IBM Canada Laboratory, Toronto, Ontario, Canada

Abstract

This position paper proposes the ***Personal Web*** as the next instantiation of the *Smart Internet Research Initiative.* Zooming in the problem scope of Smart Internet, the Personal Web focuses on the person *ME* as the center of gravity of web integration. The objective of the *Personal Web* is to empower *ME,* as a common internet user of generally limited technical skills, the autonomy and ease of control in assembling and aggregating integrate-able web elements across the web for a particular *sphere of context* of my concern. This should be so simple to do that becomes frequent everyday tasks. Typical integrate-able web elements are made available as extensions to current web sites by web domains' software programmers, to enable end users to drive and operate web integration by themselves. This user-sovereign web integration results in the generation of dynamic and high degree of personalized web artifacts for visualizations and interactions, synthesized together in a manner that is cognitively supportive; melding silos of web domains into a unified web platform to support the individual user's situational and/or persisted and repeated need for contextually relevant data and services without any programming requirement. Currently, the most popular solution towards such goal is mashups of various nuances which focuses on web sites, information and services web integration. The Personal Web proposes a people-centric integration of information, services and web content as a significantly different approach of web integration. This position paper asserts that this approach offers much better potential to realize the set goals of the Personal Web than the currently prominent mashups approach. The Personal Web is fundamentally built upon Resource Description Framework (RDF) based Linked Data of the semantic web by Tim Berners-Lee, with several significant adaptations and extensions of linked data proposed in the paper to resolve the current technical and practical limitations of mashups, in order to realize the Personal Web goals and objectives. Adaptations and extensions discussed in this paper include ***personal linked data; linked services*** and linked services scoped into ***personal linked services.*** In addition, a three-layer meta-model of the Personal Web is also proposed in order to enable more opportunities for the inference and discovery of semantic relationships between the user's sphere of context and the participating integrate-able web elements. With a defined meta-model of the Personal Web, simple user operations for web integrations can be defined and made available to users to control web integration for themselves without additional tools and widgets. I believe that the strategy set by the Personal Web will bring forth a significant step to improve users' lives by putting control in their hands in pulling together data, services and web content feeds from around the internet.

1 Introduction

Enabling general internet users, who typically have limited technical skills, the ability to pull data; services and web content together, from across the web as an integrated platform for the users' situational and/or persisted-repeated need, has been a long standing technological objective of recent web research and Web 2.0 development. There are various approaches to bring this into reality. The most recent and prominent approach towards such set goal is Mashups. Grammel et al. defines mashups as an end user driven recombination of web based data and functionality [11]. Mashups takes the approach of data, services and web presentation centric integration to solve user's need for information and services not originally anticipated by the web programmers [17]. However, constructing mashups typically involves the use of mashup tools that require deep programming skills that general internet users do not typically have, resulting in a very high barrier for adoption. To lower the barrier for adoption, widget user model for mashups is being developed recently as an alternative. Yahoo pipes and Popfly are some examples. While the widget user model claims that not a single line of code is required to be programmed, nevertheless, these mashups widgets are abstractions of necessary programmatic constructs and are not easily understood by general users unless they have prior understanding of these programming concepts such as data flow, loops etc.. Therefore general internet users still find these mashups widgets complex to understand and use [20]. Newer approach such as Mashups with programming-by-example tries to give users more direct manipulations. However, the operations in terms of what users can do with the aggregated data are very limited [17]. The limitations and perceived complexity from general web users in the various nuances of mashups are inhibitors for general internet users to perform web integration by themselves as frequent everyday tasks. Even if users manage to survive the mashup creation phase, the data and services within each participating unit of integrate-able web elements are of very low degree of personalization, making mashups inadequate in realizing the set goals of the Personal Web.

Analyzing social networking platform from the perspective as a web integration platform as an alternative to mashup provides a promising progress towards mashups' limitations that is worth noting here. For example, Facebook started with a ***social graph***, a term coined by Zuckerberg of Facebook, to capture relationships among Facebook users. Building on top of this, in April 2010, the Facebook social platform launched an Open Graph protocol as a graph for "real-world things" such as movies; restaurants etc., which are, in essence, ***integrate-able web elements*** abstracted and extended from the source web sites. Web sites can be enabled for Open Graph by defining the subset objects from source web pages and abstracted them as "real-word things" into Open Graph objects through the prescribed Open Graph protocol. By adding XML tags to the site's source HTML, these "real-word things" as Open Graph objects can become stream feeds that can be integrated into Facebook's pages such as 'wall'. Facebook provides one user operation to enable user-sovereign web integration without programming requirement, which is the "Like" operation, implemented as a "Like" button with the associated social plugins. Facebook users can use the "Like" button to create a "Like" relationship between the Open Graph object as a stream feed and the user. Facebook user can also use the "Like" button from a web page outside of Facebook, such as the sample in Figure 1. As a result, users integrate web pages themselves by using the Facebook's "Like" button, resulting in the "real-world things" represented as Open Graph objects, imbedded as subset web elements from a web page of a web site to become a part of the Facebook page. It has domino impact cascaded down through the previously established social links of the Facebook's social graph. Open Graph is implemented as Resource Description Framework in Attribute (RDFa). According to Mark Zuckerberg, this is an important step towards the creation of a more semantically aware web [24].

Figure 1 Sample Web Page with a Facebook "Like" Button

What is more significant, as per Zuckerberg's statement made in April, 2010 at the Facebook F8 Developer Conference, is that "My identity is first with Facebook and then is defined by things all over the web" [23]. This is the first closest claim towards a personal web implementation.

The notion of re-purposing social networking platform as a platform of web integration is still at its infancy stage, but with surprisingly tremendous success in terms of adoption rate. Despite of its short history, there are a few critical observations that are significant to note for its next iteration.

The first observation is that ***less is more in the meta-model of real-world things*** of integrate-able web elements. The meta-model of Facebook's Open Graph Objects and its data type is an extremely simple graph-based taxonomy. There were only a few categories of "real-world objects" such as *people; places; products and entertainment etc.,* with a very shallow taxonomy. There was only one link relationship in this graph: the "Like" link. Facebook users click on the "Like" button to specify "real-world objects" from web pages that they like. The user model cannot get simpler than this. User-sovereign web integration is achieved, yet so simple to use for general internet users that can now be done as frequent, everyday tasks. There is no programming requirement or complex mashup widgets. It drives the first massive real life adoption of an apparent semantic web site of pivotal significance. In this new paradigm of web integration, the social cascading effect because of the links established in the social graph is a critical novice element that the mashups approach cannot compare. All users from the Facebook social graph who have direct or inferred relationships with ME, as the Facebook user, are now made aware of these Open Graph objects because "I", the user, "Like" it. Cascading effect through the "Like" links can further ripple through the social network by my "friends".

On the contrary, open standards such as Web Ontology Language (OWL) are highly open and powerful in expression and rich in features and capabilities such as classes, instances with many link types. But the price of openness and expressiveness is complexity. OWL requires deep technical expertise to use and therefore results in a high barrier for adoption by everyday users. While powerful, the user does not know what to do with it, inhibiting the goal towards end-user programming [16]. When such complexity of meta-model is exposed to the user, it tremendously inhibits adoption by every day users.

The second observation is that ***defining a universal, common set of integrate-able web elements with a simple meta-model is necessary.*** It is unrealistic to expect original web pages to participate in user driven web integration with its entirety. A subset of web elements has to be abstracted from its source web site and be normalized into a set of integrate-able web elements as units of participation in user driven web integration. Such integrate-able web element in Facebook's Open Graph is Open Graph Objects. Many web domains have already followed the prescriptive steps according to the Open Graph Protocol API to abstract subset web entities from their web pages into Open Graph Objects as integrate-able web elements. There is clarity and simplicity in the prescription of such process of abstraction and normalization that existing web sites know what to do in order to contribute their sites' subset web entities into integrate-able web elements. This is also critical for quick and wide adoption.

Almost all mashup solutions do not require the participating web domains in mashups to abstract subset web entities from their source web pages in order to participate. However, all the complexity of mashing, including the abstraction and normalization from source sites, is left to be dealt by end users at the mashups creation phase. Alternatively, in the proposed personal web approach, a current web domain abstracts a subset of data, services and web contents out of its source domain in order to create integrate-able web elements to participate in the Personal Web. User can control web integration without programming requirements by using simple operations to create web integration by assembling these integrate-able web elements from across the web into a cohesive entity.

The third key observation is that ***a simple, well defined meta-model enables system generation of visualization and interactions of web integration results.*** The well defined meta-model of Open Graph Objects enables system computed web pages into a synthetic entity for the visualization and further interactions of the integration result. The Facebook "wall" with all imbedded "liked" Open Graph Objects is a good example. In the same token, the proposed basic infrastructure of the Personal Web, using RDF based Linked Data, provides much better opportunities to automate system generated web artifacts for end users to visualize and interact with the results of the user-sovereighn web integration. This is pivotal in enabling a simple user model.

The fourth major observation is that an ***RDF-based, graph oriented meta-model is designed for integration and offers much better chances for open web integration with more depth and breadth***. The RDF-based Linked Data is by nature designed for open integration. With all integrate-able web elements of the Personal Web normalized into a common RDF representation, more open and dynamic aggregation across the web is now possible.

Harvesting from these observations, the Personal Web intends to provide "ME", as a general internet user, the power and control over situational or persisted and repeatable web integration that is relevant and appropriate to MY specific requirements of content, information and services, as per my current sphere of context, so simple to use that they can become effortless, frequent everyday tasks. Like Facebook, the Personal Web proposes a people-centric integration of content, information, services. In the case of Facebook, all web integrations are done and owned on the Facebook server side. Personal Web, on the other hand, started with "ME" and my current "***sphere of context***" that orientates the web integration, integration can happen on the users' ground.

The rest of this paper is organized as followed: Section 2 covers a proposed scenario of the operations of the Personal Web. Section 3 discusses related work in the context of assessing the distance from the goal of Personal Web. Section 4 calls out the key conceptual and technical principles. Section 5 describes the proposed extensions of the RDF based Linked Data in order to form the based infrastructure for the Personal Web. Section 6 highlights the future and a set of problems and research challenges. This chapter ends off with a discussion on the potential significance of the Personal Web.

2 A Scenario of Operation

To demonstrate the Personal Web concepts and to illustrate how the Personal Web may work, a proposed scenario of operations of the Personal Web is described in this section, both from the perspective of the participating web domains and the perspective of the end users. The scenario is still a sketch and will evolve as the Personal Web continues to mature. As concepts are refined through more studies; validations of research results and user feedback from implementations, more polished version will be developed.

"Personal-Web-Enabling" an Existing Web Domain

For each web domain to participate in this user-sovereign web integration as the Personal Web prescribes, the participating web domain first needs to identify and abstract the subset web elements from the current site, such as subset units of information; subset units of services and subset web content as feeds, and normalize the abstracted subset into a set of integrate-able web elements as per the prescription defined by the meta-model and protocol of the Personal Web. When this process is completed, such web domain is considered "personal web enabled". A set of RDF based Linked Data Graphs that capture the set of integrate-able elements into a set of machine process-able form is created and made available. The resulting integrate-able web elements are to be used by the Personal Web users to drive web integrations conducted by them, integrating with integrate-able web elements from other personal-web-enabled web domains, in a manner that may not be anticipated by the originating web domains.

A Proposed Scenario of a User-Sovereign Web Integration:
In the following proposed scenario of what an user-sovereign web integration may look like, the Personal Web users articulate *a matter of concern* as a "*sphere*" of context. Around this articulated context of matter of concern, a ***Personal Web Sphere*** for that concern is formed from the finite set of normalized integrate-able web elements that come from each member-URL-s, which have been personal web enabled. The finite set of member-URL-s is defined by the user or by affirming the system's recommendation based on the computed semantic relevance of personal-web-enabled URL-s to the stated matter of concern.

For example, a user, Jo, defines a ***sphere*** for "online shopping" in order to synthesized all her online activities across the web into one integrated entity. Jo only shops from two sites online, both sites have been personal-web-enabled. After defining a personal web sphere, called "Jo's online shopping", she added two online shops' URL as member URLs for this sphere. Figure 3 shows an instance of user interface of what this may be like.

Data and Service Retrieval:
Each time, when a member URL of a set personal web sphere has been added, the corresponding member-URL's data schema subset; the user's share of data instances; and the subset of web services, and a set of web feeds from that source web domain, which have all been normalized into a set of integrate-able web elements by the source web domains as per prescription, captured in a machine process-able form in RDF based Linked Data, are loaded into and made available for the user's personal web sphere.

Data and Services Integration:
All Linked Data graphs from all member-URL-s are further integrated into one graph at the scope of the 'sphere', by relationship discovery through inferences or other means, resulting in a macro view of all integrated member as one cohesive entity.

Data and Services Cleansing and Classification
User can cleanse data or services that one does not want. User can also add semantic relationships to make the integration more relevant to him and to his situation with a mixed initiative approach.

Figure 2 A Sphere of "Online Shopping"

Figure 3 Adding personal-web enabled URLs to online shopping sphere

Data and Services Visualization and Interactions
Visualization and web interfaces of the merged RDF graphs from all member sites are be generated dynamically. There has been research work done in visualizing generated RDF graphs such as Haystack [13]. This allows the user to view and operate on the integrated data and services as one aggregated entity.

This proposed scenario of operations is a bookmark-like user model in the creation of personal web space as a way of user controlled web integration. While mashups has been viewed as a form end user programming of the web to enable web integration, the Personal Web proposed to take out the concept of programming from the control of general internet users in web integration and maintain the objective of user sovereign web integration with no programming requirement. It is asserted that this is significantly simpler than the mashup web integration alternative which typically requires additional mashup tool or mashup widgets.

3 Related Work

In the last five years, there has been rapid development in technologies towards the goal of enabling end users to aggregate web content; data and functions from multiple websites into an integrated unit, for a given purpose or to solve a specific problem. What makes this exciting and inspiring is the boundless innovative possibilities when the power and control over the web is put in the hands of end users, to create what may not be previously possible and anticipated. The survey and discussion in this section attempts to assess to what extent has such stated objective a reality for general non-technical internet users, to the point that they can perform such web integration with very minimum effort as frequent everyday tasks without any programming requirement.

2.1 Mashups

Mashups are the most popular tools of such enablement. Mashup is generally referred to as a web application that integrates data and sources from multiple sources to provide a unique service [20]. Studies have been conducted on the different kinds mashups designed for different functions and purposes. As a result, a set of abstracted mashup patterns have been identified [22]. In surveying the various models of mashups available today, Figure 4 shows the relationship between end user's skill requirement and the extension requirement of the participating web domains. While the potential of mashups to help users to solve their everyday problems is great, yet to this day, real world adoption for end user web integration as everyday activities is still not a reality.

Figure 4:

Progression of End User Enablement in Web Integration

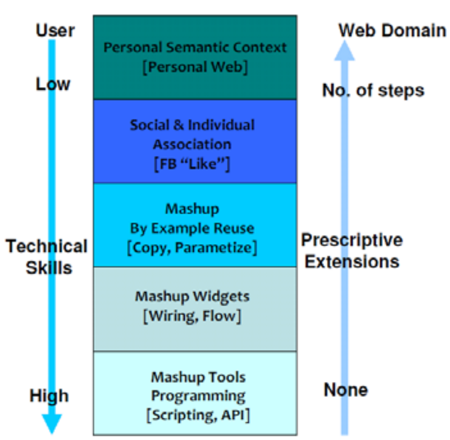

The early offerings of mashups often require technical and programming skills beyond what one could expect from general internet users. For example, for Mash Maker to extract data from multiple web sites to be aggregated, an extractor has to be first created for such aggregation. The creation of extractor requires some programming skills [16]. Microsoft Popfly's involves a visual dataflow language. Google Mashup Editor uses extended HTML and Javascript heavily [11]. All these illustrate a high barrier for general internet users to freely and frequently construct data aggregation across web sites as their everyday activity.

The widget approach of mashup intends to eliminate the need of coding. Yahoo's pipes, Popfly are examples of this widget approach. Mashups widgets represent different web sites. Widgets are to be connected together through wiring. But such approach, while an improvement over the programming required for Mashup development tools, they have not yet able to offer satisfactory alternative. The number of existing widgets is usually just too many to be manageable and they are difficult to find. While it is true that with mashup widgets, users are not require to code or write scripts, however, mashup widgets are still too difficult for general internet users unless they have some understanding of basic programming concepts (such as loops; data types etc.).

Mashup by example is the latest development. It attempts to lower the deep technical skills required in the widget approach of mashup. System such as Karma enables its users to extract data from a web site into a data table through demonstration, and uses XPath generation to find similar data and copies into the table as well. However, such integration is only limited to data integration and is good only for ad hoc integration that does not require persistence. The requirement of service integration is not being addressed. System such as Vegemite is implemented spreadsheet like user interface to provide users direct manipulation [17], but has the same limitations in persistence and lack of support for service integration.

One can conclude that to this day, there has not been one single mashup model that is made available to general users, easy enough to be used for frequent, every day web integration activities, that can support robust and deep integration of data, services and web content feeds, all within one mashup model with high degree of personalization.

Facebook Open Graph

Recent research has proposed to leverage technologies established in semantic web to create solution that can relax the current limitations of mashup previously described [3]. The Facebook instance provides an interesting case study of using a social networking platform also as a web integration platform. Facebook uses RDFa based "Open Graph Protocol". The Open Graph Protocol prescribes how a given web domain should "tag" the HTML of their site in order to enable the web page with "Like" button that their user can select. The site can also follow the Open Graph Object tags to create stream feeds with a "Like" button to participate as an imbedded object of face pages. When a web domain has followed the steps of prescriptive extensions of their sites according to what the Facebook Open Graph APIs have prescribed, Facebook provides its users a simple point-and-click user model to establish relationships between "real world things" and people of the social graph and network. With the connections established in the Facebook's social graph, a cascading effect of social impact is generated with the thing that one given Facebook user "Likes". The degree of technical skills required from the user of this web integration approach is very low.

Facebook is an interesting experimentation of merging web 2.0 with semantic web and it is not without its limitations and shortcomings. For examples, the available categories of Open Graph Objects are very limited in terms of comprehensiveness in representation of real-life objects. Also, the current meta-model of Open Graph Objects is limited to define a fragment of web feed by listing the set of references of images; text descriptions etc. from current web pages. The Open Graph meta-model falls short in defining objects for services and backend data such as data from relational databases to participate in this user-sovereign web integration. Scaling up, the requirement of security, such as authentication and fine grain authorization of Open Graph Objects and issues of privacy must also be addressed. At its infancy, critics of Facebook Open Graph are quick to point out that "Open" is not open because all aggregations have to be conducted by users of Facebook and must place on a Facebook server. While the simplicity of user operation of the "Like" button drives fast adoption, there are more relationships in the real world than the one and only "Like" relationship.

It is also important to point out that the mashups approach focuses on aggregation of integrate-able entities such as data, services and web feeds as different mashups types. For the longest time, technologists have attempted to define ontology that best models the *relationships between entity and entity*. The social networking approach, however, focuses on simple *abstraction of relationship between "a person" and his "integrate-able entities"*. The "Like" relationship is an answer to the question: how a user typically relates to this category of "things" to fulfill the purpose and reason of the user's web integration. In the case of Facebook, one of the purposes of web integration is to create an integrated page of all things that the user likes across the web. This realization provides invaluable insight in the consideration of adapting the semantic modeling to a model of *relationship between people and entities* instead of just focus on relationship modeling between entities and entities. All these lessons learned help to shape the Personal Web to realize its set strategy and goal.

4 The Personal Web Principles

This section attempts to capture an initial set of characteristics core to the Personal Web, both at the conceptual and technical perspectives. These principles are to be validated and adapted as the research and development of the Personal Web evolves.

4.1 The Conceptual Principles of the Personal Web

#1. The user "I" is the center of gravity for web integration

Facebook's web integration is oriented around the users of the social graph, for the purpose of aggregating social impact of all real-world objects liked by an user as an object of the social graph, cascading through the established "a friend of" relation through out the social graph. The Personal Web, on the other hand, orientates web integration around the user as the center of gravity for web integration. All data about "ME" across the web, such as all user IDs and passwords, all account numbers across the web that belongs to ME such as frequent flyer numbers, can be aggregated into an integrated data entity represented by one consolidated Linked Data GraphThis "ME" object with the practical benefit of sparing ME as an user, from repeatedly entering static data entries such as user ID and password; account information etc., every time, resulting more pleasant user interactions and experience [19].

#2. My Context Shapes the Scope and Semantics of Web Integration

Dey defined context as any information that can be used to characterize the situation of a person, place or object, in order to bring the most relevant content and services, for the user's benefit of the situation [9]. The Personal Web performs web integration around a user, and the user's personal semantic context, called "Personal Web Sphere". A little semantic goes a long way [14] for the discovery of semantically relevant data; services and web feeds for a specific context that the user specified.

#3. The user "I" is completely sovereign in controlling the web integration

Building on the Web 2.0 assertion of the Web is a platform, this principle of personal web takes a step further to see the web as one concerted platform to support MY every day situational or persistent and repeated need of data and services, that I, as the general non-technical users of the internet can do as frequent, everyday tasks. Mashups is sometimes being referenced as end user programming of the Web. The programming concept is to be removed from the Personal Web and instead, it asserts that the user model of the Personal Web has to be so simple that frequent everyday web integration totally operated and controlled by general internet end users is made possible without any programming requirement.

#4. Let the Web works on my behalf

The cognitive load reduction is one of the major goals of the Smart Internet. More advanced features of the Personal Web will be developed to include these concepts such as enabling an individual user to persist some previously persisted form of web integration as scheduled tasks or events as prospective memory tasks; or to set a certain persistent and repeated task sequence in automated into autopilot mode, running in the web platform as a batch processes with the user plays the role of a supervisory controller [1].

#5. It has to be a social web after all

It is important not to lose sight of the important fact that the special user "I" is also a part of the social network. While there is no concrete validation of user requirement, but hypothetically, sharing my personal web spheres with my friends on the social network, or my friends invite me to participate in his personal web sphere are some possible, futuristic scenario. Asking my friends what URL-s are most revelant to my Sphere is another practical scenario to add. The domino impact of the social network in influencing the adaptation of the most optimal and effective web integration for a given semantic context for a particular profile of users is too important to ignore.

4.2 The Technical Principles of the Personal Web

This section calls out the technical characteristics of the Personal Web architecture and design. These technical principles are intended to drive design and implementation decisions. Abiding to these technical principles in all components of Personal Web implementation is pivotal to its real world adoption.

#1. The Programming Goal of the Personal Web is to abstract out the programming complexity for its users so as to Pass Control from the programmers to the End Users

The goal of any API protocol and programming model of the Personal Web is to produce intermediary web artifacts, usually in machine readable and process-able form, for the sake of enabling the Personal Web users the control to aggregate these parts into something of a bigger context as the user so desires. The design goal of the Personal Web components and intermediary artifacts is to produce independent units of web composition which are open for integration by design of the software engineers who produce them. By putting these units of composition in the hands of general users, with a simple set of user operations defined for user sovereign web integration, end users have actual control over how the web as a platform can work for them, which is not possible with today's monolithic style of web applications.

2. Less is More in meta-model for semantics and user Operations for web integration

This principle can be applied in different components of the architecture of the Personal Web. For example, general internet users cannot handle the many great features of Web Ontology Web (OWL) to express relationships. Even high skill technical users need some time to compute if a given OWL instance as an 'individual' of a certain member class is "equivalent to" another class but "disjoint" with some other classes. This is a drastic contrast to Facebook's one and only relationship of "Like", from which simple user operation of the "Like" button can be derived that general users with limited technical skills can perform with minimum effort.

#3. Design for Open Integration

It is also important to abide by the principle of openness and standards. The intermediary of web artifacts of the Personal Web is designed and architected to be a part of something else that was not originally known and anticipated. This critical requirement of openness for integration is the major reason that the RDF Linked Data is the selected candidate of technology infrastructure instead of the mashup widgets of many nuances as the foundation of choice for internal representation.

#4. Model NOT relationship among things, but model people's relationship with things

Current usage of ontology seems to focus on the modeling accuracy of concepts and objects relationships in representing knowledge of the real world. Ontology language such as OWL has been, for decades, focused on modeling the real world things and knowledge with high degree of accuracy but completely lost sight of the importance of ***why*** and ***how*** people relate to things in real life. The "Like" link of Facebook is profound in the sense that it models how people relate to things in real life according to the user's reasons of web integration, instead of how "things" relate to each other. For example, while it is true in knowledge representation that shop carte "is equivalent to" shop bag, it is more important to assert that I, the user, can "join" them as one semantically for my purpose of seeing all my shopping entries online as an integrated entity. In this example, the “Join” relationship that models the purpose of my web integration has more significance than the knowledge of shop bag and shop carte are equivalent in semantics. More research needs to be done in this area in people-and-entity relationships in order to define well-abstracted and useful user operations for web integration.

5 The Basic Enabling Infrastructure of the Personal Web

The basic enabling infrastructure of the Personal Web is ***linked data***, whose objective is to make the Web a “***Global Data Space***” by enabling data published from diverse sources of web domains to be connected with formally typed connections, in order to make the web a more intelligent source of information for its users [7]. The original goal of linked data is to “stitch together the world’s structured information and knowledge to answer semantically rich queries” [4].

Linked Data is built upon ***Resource Description Framework (RDF),*** developed with the motivation to provide an open information model that can interwork and combine data from multiple disjoint sources in order to derive new information that can be processed automatically at the internet scale [21]. The underlying structure of any expression in RDF is a collection of triples. Each RDF triple consists of a ***subject, a predicate and an object.*** Each triple represents a statement of a relationship, which can be represented by a directed graph. The meaning of a given RDF graph is the conjunction of all assertions, represented by all triples it contains. As a result, it provides a generic, graph-based data model to describe things in the world [15]. Linked Data builds on top of RDF by defining formal ***typed links*** as connections between data from diverse domains using the construct of ***predicate*** from RDF in order to link arbitrary things in the world to create “the Web of Data” [7]. The subject, object and predicate of the RDF triple are all URI-s. Each identifies a resource or a string. Predicate’s URI specifies how the subject and object are related.

The notion of developing personal information systems consisting of focused subsets of information, highly relevant to a particular user in order to deal with information overload of the user from the web platform, dates back at least as far as the personal web space vision of Abrams et al [1] in 1998. While Abrams et al focused on personal bookmarking, others looked for enhanced forms of search, but still within a linked data approach. However, the current view of the Personal Web goes beyond that earlier linked data approach to consider linked services that support not just information exploration and management, but many other tasks as well.

This paper proposes to extend the ***typed links*** of ***Linked Data*** in order to a build an infrastructure that instantiates the strategy of the Personal Web into reality. The adaptations and extensions proposed can be summarized into the following parts: #1 Re-purpose the global linked data into a personal scope of linked data called ***the Personal Linked Data***. #2-a. Extend the participating object nodes of the RDF to include Restful services that this paper terms ***Linked Services***. #2-b. Re-purpose the global scope linked services into the ***Personal Linked Services.*** #3. Integrate the *personal linked data* and the *personal linked ser*vices into one RDF linked data graph as one ***Personal Web Link***. #4. Provide an initial ***meta model*** to define and normalize these integrate-able web entities of the Personal Web into an open, universal model, prescriptive for current web domains to follow in order to extend the current site to take part in the Personal Web.

5.1 The Personal Linked Data

It is important to note that the Personal Web has a different scope and objective than that of the original Linked Data. The original concept of Linked Data is to connect local data for the global web. Personal Linked Data of the Personal Web connects user’s global data for and about the user, as an individual. The vision, goal and objective of the Personal Web, from a pure data perspective, is a ***gathering of global data about me, for me and relevant to me across the global web into a synthetic, integrated data entity of ME across the global web.*** We called this the “***Personal Linked Data***”.
From the scope and perspective of a given web domain, the web domain’s *Personal Linked Data* for a given user is the proper subset of data entities that belongs to the user and/or public data subset relevant to the user, including structured and unstructured data, as well as web content feed, from the source web site, normalized into one RDF Linked Data graph in representation. For a web domain to be personal-web-enabled, the web domain needs to define the mapping for a given user’s proper subset of data from the web domain. All user’s proper subset of data, including structured, unstructured and web content feeds that the web domain owner defined for the users will be normalized into a linked data RDF graph for that user.

From the user’s perspective, the user can gather Personal Linked Data RDF graphs from several web domains and group them together into one semantic context. In this way, multiple Personal Linked Data RDF graphs can further discover and infer relationships and links across the graphs, and also how they relate to the semantic context, which captures the user’s purpose and intent. Joining these multiple Personal Linked Data RDF graphs, from across web domains but semantically grouped within a set context, enable simple but user driven web integration.

5.2 Linked Services

Another extension from the original Linked Data proposed; is to extend the original concept of linked data beyond the scope of data and information but transposed to and applied in the scope of ***web services*** *as units of web functions.* In other words, the ***subjects; objects and predicate*** of the RDF triple and its typed links are used to express Restful service. Such extension of the Linked Data for the representation of Restful web services is being termed as ***"Linked Services"*** in this paper. Fielding's architectural design of restful services [10] in 2000 constrained each service as an URL, making restful services fit architecturally to be integrated and to participate in the linked data graphs of the Personal Web. Recent research has made progress in defining different methodologies to normalize different kinds of web services into Restful services, with the representational URI as resources. As a result, "services" as user tasks, once being normalized into restful services, can be transformed into important participants of the Personal Web. Theoretically, like its linked data counter part, "linked services" is default to the global web scope.

5.3 The Personal Linked Services

In reality, having a global scope of services is not practical or realistic. This paper asserts that services are often invoked and executed by and for a web user within a very well defined context; and the therefore, there is a much higher value proposition when the concept of "***linked services***" is applied to an individual personal web scope of ME as a web user instead of to the global scope of services. Such transformation is being termed "***personal linked services***".

From the scope and perspective of a given web domain, the web domain's *Personal Linked Services* for a given user is a proper subset of services entities that the web domain wants to make available to its users as integrate-able units of functions and tasks, normalized into one RDF Linked Services graph in representation, according to the meta-model of the Personal Web. What makes the linked services personal? Not only the choice and combination of services can be individualized, but the input resources, because the user is known, can be pre-set as per the identity and the data of the set user.

From the user's perspective, the user can gather Personal Linked Services RDF graphs from several web domains and group them together into one semantic context. There is great potential for future work to explore how to enable, and what user operations of the user model for the user to express the sequencing of these services across web domains, resulting in an extremely simple, user-sovereign process management web tool without extra tool widgets.

5.4 The Personal Web Sphere

The semantic context, which represents the reason and purpose for a user's need to perform web integration, driven and control by end user, is called a "sphere", as a sphere of context for integration. In other word, the "***sphere***" is the center of gravity that the user creates, as a unique individual, as the gravitational factors for web integration. The semantic links and relationships related to the 'sphere' drive the inference and discovery of related integrate-able web elements, resulting in a much deeper integration in web resources that yield a high degree of personalization that common user can handle as frequent, everyday tasks.

5.5 The Three Layered Meta Model

It is important to establish a meta-model for the Personal Web, upon which a universal and open framework can be developed. This is pivotal in enabling current web domains to extend their current web sites in a prescriptive manner in order for them to take part in the Personal Web. With a model to normalize to, end users can now operate with integrate-able web elements which are designed to be a part of something that may or may not be previously anticipated. In addition, with this defined meta-model, automatic generations of web artifacts for user visualization and interaction of the integration outcome can be developed. Automatic discovery of links and relationships through inferences and other means are then made possible.

This paper calls out three primary categories of integrate-able web elements for the Personal Web as participating units in this user-sovereign web integration, designed to be open and integrate-able with each other. They are, namely, (i) units of data, both structured and unstructured; (ii) units of web content in the form of feeds, (iii) units of services in terms of functions and capability. All three categories are represented as RDF Linked data objects in its machine process-able form. The strategy is to define one set of simple universal user operations for the Personal Web users to synthesize instances of all categories of web elements into something novice as a new web integration approach that end users can do without any programming requirements.

Facebook has a preliminary deployment with its Object Graph Objects, which in essence are units of web content in the form of feeds, being imbedded with other objects as an integrated Facebook page. The Personal Web proposes to add the categories of units of data and units of services to take this web integration approach to the next level of maturity. In addition, a simple system to associate semantics and meaning to these integrate-able web elements is also proposed, with such, opportunities and accuracy in links and relationships discovery within and across RDF based linked data graphs will be greatly increased. A mixed initiative approach is suggested for such extension. Such web integration infrastructure has high potential to result in an increased level of depth, breath and openness of web integration, with a high degree of personalization, and a low barrier to general users that its mashups approach cannot compare.

This paper proposes a three-layer Personal Web meta-model by leveraging Chen's Entity Relationship (ER) model that was originally intended as a top-down database design tool. The ER model's three-levels of progressive perspectives of data, (namely, the "entity" level as the top conceptual level of abstraction; the "data structure" level as the concrete level of the previous layer, and the "data record" level as data instance level) is being adopted [8]. The Personal Web's ***Concept*** layer; ***Model*** layer and ***Instance*** layer come from such adoption (See Figure 5), with RDF nodes in each layer properly typed as ***concept-nodes; model-nodes and instance-nodes***. The relationships and links between RDF object nodes within each layer and across layers are expressed as RDF predicates.

5.5.1 The Model Layer and the Instance Layer

#1. Web Content as Units of Integrate-able Object Nodes

The Data models and schema of various units of web content should be captured as captured as ***web-content model object node***. The XML schema for the Open Graph Object is an example. Many instances of such Facebook feeds were created by the participating web domains, and should be captured as ***web-content instance object nodes*** within the meta-model system of the Personal Web.

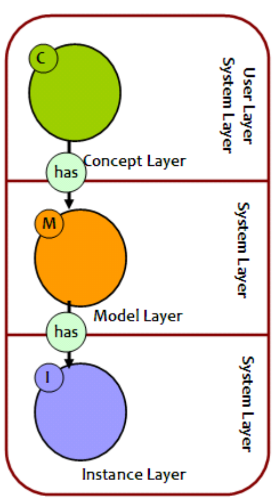

Figure 5 The Personal Web's Three-Layer Meta Model

#2. Data as Units of Integrate-able Object Nodes

The structured data schema from the source web domain (e.g. the data schema of a shopping carte from www.shopx.com) should be captured as ***data model nodes*** in the model layer. The actual data instances (e.g. the actual shop carte data records of purchases by user "Jo" from shopping site www.shopx.com) are captured as ***data instance object nodes***. All levels of traceability are maintained.

Future work needs to analyze and study real life interactions between people with these object models and instances in order to discover and abstract the most engaging and useful people-object relationships. Such findings will provide great insights for the definition of new link-types for simple user operations of creating purposeful web integration with no programming requirement. This paper asserts that these people-object relationships are ***real-world verbs*** (such as "Like" in Facebook) and can be represented as RDF-predicates**.**

#3. Services as Units of Integrate-able Object Nodes

The service meta-models should be captured as service models in the model layer. Two important service models should be called out, namely, the ***synchronous and asynchronous services***. The corresponding actual deployed service instances are captured in ***service instance object nodes***.

Future research work will refine and validate these service models for synchronous and asynchronous services. The initial sketch of a service model, kept in the model layer, for synchronous services includes a *resource node* and a *task node* with ***create, update, read and delete*** as verb-like predicates. The initial thought for an asynchronous service model is contains a *resource node* with a *task node* to be linked with ***publish and subscribe*** verb-like predicates. More real life deployment is needed to validate such service model in RDF representation.

5.5.2 The Concept Layer

Unlike Chen's ER model, the concept layer functions more as an ontology or taxonomy layer than the data entity abstraction of the original ER design [8]. Its main mission is to enable the linking of multiple model nuances in the model layer with one or more instances of a concept node in the concept layer. Since traceability is maintained across all three meta-model layers, algorithms can be designed to synthesize a unified data from disjoined silo web sources into integrated operations (e.g. "checkout" all shop carte for all member URL-s of my shopping sphere). Nuances of similar concepts can also be linked at the concept layer as well.

Ontology and data integration problems are well known issues. Each source web domain has its local conceptualization, ontology and 'tagging' system being handcrafted by database administrators and domain knowledge engineers. They are disconnected and disjoint from other web domains [12]. It is important to look into how semantic heterogeneity is being supported to achieve data integration practical to the users [5]. General user should also be given an opportunity to tag and map these concepts for one's personal matter of concern [18]. A mixed initiative approach in which both system automation and user participation is being proposed. Web domain may create and link some concept nodes with the corresponding data model object nodes to enhance user's web integration.

At this early stage, it is important to call out the concept node of "ME" and "SPHERE" are predefined and are special concept nodes for web integration. Future work is to be done in this and to be discussed with more detail.

Concept Nodes for the Real World Modeling: People, Places, Events, Tasks and Things

Modeling relationships between "things" only answers the question of how real world objects are related to each other. But it does not answer the question of how "people" relate to "objects" or "things. By carefully analyzing and abstracting relationships between people and things, and empirically answering the question of how/why "people" use/relate to "things", there is a better chance that more purposeful web integration for user's problem at hand can be enabled.

When people relate to things centers around a given context, the temporal and location dimensions are essential aspects. In this light, concept nodes are classified into concept nodes of the following category: People, Places, Events, Tasks and Things, Future work is needed to validate such taxonomy of concept nodes.

Types of Predicates as real world verbs

With what has been previously established, the RDF triple of subject-predicate-object can be applied to represent the ***people-action-thing pattern*** in answering the question of how/why people relate to things. The "Like" link of the Facebook Open Object Graph Object is a great example of such pattern. The predicates of the concept nodes are verbs as action words in modeling real world concepts.

6. Discussion and Future Work

This position paper is to set the vision, strategies and goals of the Personal Web and proposes to use RDF Linked Data as its technology base. Unlike the original strategy of Linked Data, which connects local source data in the global web, the Personal Web re-orients web integration and connects the global web data for and about ME in my personal web sphere to synthesize ME in the global web. Personal-web-enabled sites provide integrate-able web elements in a machine readable form of a RDF graph, with simple user operations that requires low technical skill level, that "I" the user can link with other RDF graphs from other silos web domains as a form of web integration in a manner that is distinctively different from the prominent mashups approach. While such novice web integration approach has great potential, there are lots of remaining challenges and open research issues related to the enabling of the Personal Web. The major open research areas are pointed out in the following. It is not claimed to be comprehensive.

6.1 Personal Linked Data

There are many open research challenges related to the concept of Personal Linked Data. The first challenge is to define a practical framework in order for existing web domains to be able to identify and capture the data model subset mapping of structured and unstructured data that will be of high value for its users' personal web sphere integration. The second challenge is that even if such data model subset is identified, what approach should be taken to identify the data instances subset mapping that captures only the data instances that belongs to a specific user. The third challenge is how to normalize the identified structured and unstructured data into RDF and yet preserves the original semantics that can promote links discovery. DBpedia has established a framework of such transformation [6]. However, the normalization loses some degree of items independence, resulting in the loss in semantics and opportunities to discover new relationships.

6.2 Linked Services and Personal Linked Services

Research in the relationship of RESTful services and Linked Data is still very new. Recent research proposes to describe RESTful resources as semantic resources as in RDF Linked Data, with the strategy to harvest the data that's already out there [2]. The strategy of Linked Services is to use RDF Linked Data construct to express the service models of synchronous and asynchronous RESTful services in a form that they are normalized into units of integrate-able web elements that end users can control.

This creates several interesting research challenges. First of all, independent from the goals of the Personal Web, normalizing web services of all models into one common RESTful services model is by itself a significant challenge. Secondly, expressing RESTful services models in RDF Linked Data Graph, that can become unites of web integration such that end users can operate with, needs careful analysis and proper abstraction. However, if significant progress is made, user can control without programming requirements process flow such as sequences of services call, or to create script-based like agents for auto-pilot mode execution. How to make a well represented RESTful service in RDF Linked Data personal? This should be studied, not only to cover individual user's personal choice in the list and combination of services, it should also cover how an service instances can be preloaded with the user's personal data in order to improve the simplicity of service calls.

6.3 Lightweight but Expressive Semantic Infrastructure

Learning from the Facebook deployment example, an extremely simple data model and semantic infrastructure is pivotal in real world adoption. A lightweight but expressive semantic infrastructure is critical to the Personal Web in order to build associations between concepts, such as online banking, with the integrate-able web entities, including web feeds; data; and services.

One possible technology choice is Bergamaschi's "Common Thesaurus", constructed by analyzing Object Descriptive Language (ODL) and generalized into a very simple relationship model. There are three types of relationships in Common Thesaurus: "Synonym-of" and "Related-term" are the two symmetric relationships. "Broader-term" is the non-symmetrical relationship [5]. Adding mapping rules and attributes to this simple relationship model enables query and visualization of data source across different web domains. Other work done in this area is ConceptNet, It has a more elaborate relationship model built on WordNet with about twenty relationships [18]. More studies need to be done to define a practical semantic infrastructure that is simple to use, like the Common Thesaurus, yet expressive, like the ConceptNet.

6.4 Visualization and Interaction of User Sovereign Web Integration

Visualization of a complex integrated RDF graph by applying visualization techniques that is semantically accurate but slick in the eyes of the user is a key to general adoption and is an important area of research. Developing a framework to parse machine readable RDF and transform a graph based representation into a user friendly visualization artifacts that users can interact with is an important problem to solve. Future research work cannot ignore the studying of adaptability based on factors of the context such as location, device type.

6.5 Practical Issues

It is also important not to lose sight of the non-technical but practical issues such as the trust of the user: will general user trust a Web integrated ME to reside on the server side or worrying about losing identity if it lives on the user's device? There are also privacy considerations. How much identity of me could be stolen if I lose this integrated ME object? How much trust factor in the user to depend on the generated integration results? It is desirable to have a life-deployed instance of the Personal Web to provide some empirical, real life data of usage and feedback to make progress.

7 Conclusion

The Personal Web intends to provide "ME", as a general internet user, the power and control over situational or persisted/repeated web integration that is relevant and appropriate to the user's concern of the moment. The set goal is that it should be so simple to do that become general internet users' frequent, everyday tasks. We all know that this is not a reality yet. One has to conclude that the barrier of adoption is high due to its complexity. An entirely different paradigm of a social network oriented approach in web integration, using Facebook as an example, has great adoption with simplicity in user model. This paper proposes to learn from the case study of Facebook in order to propose yet a newer paradigm of web integration. Instead of a web integration oriented primarily around a social network, the Personal Web proposes a paradigm of web integration primarily orientated around ME as the user and my semantic context as a "Sphere" that scopes the integration gravity. While ME as a user is also a part of the social network, it comes second in priority after ME as the web integration driver. The Personal Web technology infrastructure is based on RDF Linked Data. The openness of such infrastructure offers tremendous opportunities with promising potential that this proposed strategy would result in web integration with more depth, breadth and degree personalization that users can do as a frequent every day tasks that the mashup approach cannot compare. Mashups today can be categorized into three typical categories: web content feed, information and services. There is no mashup today that integrate all three categories of web integration in one cohesive system to be put in the user's hand. The Personal Web claims to provide all three categories integratedly, which current mashup approaches cannot compare..

Acknowledgments

I would like to give my special thanks to Dr. Mark Chignell, Dr. James R. Cordy and Dr. Yelena Yesha for their great contributions in many ways to evolve Smart Internet into its next instantiation of the Personal Web. Your dedication and enthusiasm in pushing new ground in these new and exciting areas of research is energizing and inspiring. I would also like to thank Anatol Kark and Emilia Tung for their tremendous support. I thank the CAS Technology Incubation Lab. team members and CAS Research Staff Members, Jimmy Lo and Alex Lau, for their technical input and excellence in this project.

About the Author

Joanna Ng is currently the Head of Research at IBM Canada Software Laboratories, Center for Advanced Studies. She is also a Senior Technical Staff Member of IBM Software Group. She has held various senior management and architect positions in product development and software strategy within IBM. Joanna is an IBM Master Inventor with a long track record of profitable innovations. She has been granted over twenty five patents from various countries in research areas related to Smart Internet such as mobile commerce, voice-enabled portal, commerce portal, retail industry solutions; service-oriented architecture (SOA); asset repository; semantic and web technologies.

References

[1] Abrams, D., Baecker, R. and Chignell, M., "Information Archiving with Bookmarks: Personal Web Space Construction and Organization", CHI Conference, 1998.

[2] Alarcon, R. and Wilde, E., "Linking Data from RESTful Services", LDOW, April, 2010.

[3] Ankolekar, A., Krotzsch, M., Tran, T., Vrandecic, D., "The Two Cultures", ACM WWW 2007.

[4] Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z. "A Nucleus of a Web of Open Data", 2007.

[5] Bergamaschi, S., Castano, S., Vincini, M., "Semantic Integration of Semistructured and Structured Data Sources", ACM Sigmond Record, 1999.

[6] Bizer, C., "DBpedia – a Crystallization Point for the Web of Data" Web Semantics: Science, Services and Agents on the World Wide Web, Volume 7, Issue 3 p. 154-165, September 2009.

[7] Bizer, C., Heath, T., Berners-Lee, T., "Linked Data: The Story So Far", International Journal on Semantic Web and Information Systems (IJSWIS), 2009.

[8]Chen, P.P.S., "The Entity-Relationship Model – Toward a Unified View of Data", ACM Transactions on Database Systems, Vol. 1, No. 1, p. 9-36, March 1976.

[9] Dey, A.K., "Understanding and using context", Personal and Ubiquitous Computing, 5(1), 2001, pp. 4-7.

[10] Fielding, R.T., "Architectural Styles and the Design of Network-based Software Architectures", Doctoral dissertation, University of California, Irvine, 2000.

[11] Grammel, L., Storey, M. , "A Survey of Mashup Development Environments", University of Victory, 2010.

[12] Guarino, N., "Formal Ontology and Information Systems", Proceedings of FOIS'98, Trento, Italy, IOS Press, pp. 3-15, June 1998.

[13] Huynh, D., Karger, D., Quan, D, "A Platform for creating, organizing and visualizing information using RDF", Semantic Web Workshop, 2002.

[14] Hendler, J., Keynote of International Semantic Web Conference, 2003

[15] Klyne, G., Carroll, J.J., McBride, B., "Resource Description Framework (RDF): Concepts and Abstract Syntax", W3C Recommendation, Retrieved http://www.w3.org/TR/rdfconcepts/

[16] Ko A.J., Myers, B.A., Aung, H.H., "Six Learning Barriers in End-User Programming Systems", VLHCC'04: Proceedings of the 2004 IEEE Symposium on Visual Languages – Human Centric Computing, pages 199-206, IEEE Computer Society, 2004.

[17] Lin, J., Wong, J., Nichols, J., Cypher, A., Lau, T.A. "End-User Programming of Mashups with Vegemite", ACM IUI'09.

[18] Liu, H., Singh, P., "ConceptNet: a Practical Commonsense Reasoning Tool-kit", BT Technology Journal – Vol 22 No 4, October, 2004.

[19] Ng, J., Nigul, L., Litani, E, Lau, D., "End user controlled web interaction flow using service oriented architecture model", The 2nd IEEE Workshop on Enabling the Future SOA Internet (EFSOW), IEEE Globecom, 2008.

[20] Tuchinda, R., Szekely, P., Knoblock, C.A., "Building Mashups By Example", ACM IUI'08.

[21] W3C, "Resource Description Framework (RDF): Concepts and Abstract Syntax, W3C Recommendation, 10 February, 2004.

[22] Wong, J., Hong, J. "What Do We 'Mashup' When We Make Mashups?", ACM WEUSE IV'08.

[23] Zuckerberg, M., Facebook F8 Developer Conference, April 2010.

[24] Zuckerberg, M., Taylor, B. "The Open Graph Protocol", www.opengraphprotocol.org, April 2010.

Modeling and Analysis of Personal Web Applications: A Vision

Marsha Chechik, Jocelyn Simmonds, Shoham Ben-David, Shiva Nejati, Mehrdad Sabetzadeh, and Rick Salay

Shiva and Mehrdad are with Simula Research Lab, Norway. Other authors are with Department of Computer Science, University of Toronto

1 Introduction and Assumptions

In this paper, we attempt to identify our vision of what personal web is. We then provide a challenge problem for such a vision and discuss our assumptions of how other fields of computer science can contribute to executing this vision. We then discuss where the "traditional" software engineering tasks of specification, modeling, monitoring and verification fit into the vision.

What is Web 3.0? We share the vision of the workshop organizers that it is focused on an individual user (well, consumer) and is trying to elicit and execute this person's goals, through preferred information collection devices and with cooperation with trusted individuals. For example, traditional web applications such as commerce and banking offer a particular interaction with the user and his/her data. Data is stored in the database on a particular application (e.g., shopping list or wish list), and the user is being offered a particular workflow that determines the interaction of the user with the system (e.g., on amazon.com, such things include looking for something, doing a price comparison, determining a particular vendor to go with, choosing the type of shipment and the payment method).

Instead, as users, we find those parts of the different applications which are useful to us, and then combine them in ad-hoc ways. For example, when buying electronics, we (a group of Toronto-based academics) first check amazon.com to look at the models and reviews. Amazon.ca has a much smaller product selection, and very likely will not carry the desired product. Instead, we look for the equivalent models on other Canadian retail sites. After comparing prices and shipping options, we make a purchase. In other words, we have an informal workflow for buying electronics online. On a good day (no paper deadlines, no screaming children, no advisors interrupting with urgent requests), we have time to do each step of this process meticulously and end up completely satisfied with the end result. But when pressed for time, we start skipping steps (since nothing is automated), e.g., we will order directly from amazon.com and pay extra shipping and customs charges, or shop our favorite Canadian retailer instead of comparing prices, and sometimes miss out on a sale.

Our vision of personal web is that it allows users to define and manipulate personal workflows, populated by their favorite vendors and information sites. For example, users would download the "web for a Canadian shopper" workflow, that implements the above workflow for shopping for electronics (clearly, it is right about the same no matter what is being purchased) and then modify it to suit their goals. Other clear workflows are for organizing a dinner and a movie outing (choosing an interesting movie - and checking appropriate sites to determine what is good - a time that works and at a location which is reasonable to get to and that has a restaurant close by that the person executing the workflow would like to visit - coordination with the person's date and the restaurant review sites).

Thus, our definition of a personal web is *Support for identifying and executing (and monitoring and fixing) "mental" orchestrations, or workflows, for incorporating multiple services in order to accomplish complex personal goals*.

The problem is clearly non-trivial and its solution requires a collaboration from various areas of computer science. Here is a very partial listing of the issues:

Data storage. Personal web needs an ability to store the state of each user within their workflow as well as the associated collected data, since it is no longer done at the vendor's site. We think that cloud computing can readily provide a solution for this challenge.

Turning the web into services. The workflow-based vision means the ability to invoke services rather than browse the web. That is, the web should be turned into a collection of such services. We think that the semantic web research community has a lot to offer on this topic.

Services specification. It is essential to have some notion of specification for services, at least to determine whether a particular service can be invoked at a particular step of the workflow but of course to also discover services (don't you want to know that a local computer shop is having a sale and it might be cheaper to go there rather than continuing the on-line shopping experience?). Personal web is not unique in this challenge - it is essential for creating quality web service applications under existing technologies.

Architectural support. Given that users define these personalized workflows, where are they being executed? How are services being "strung" together? Again, there seems to be a lot of success in existing technologies for creating mashups, yahoo pipes, etc.

Modeling and analysis. This is the purpose of this paper - trying to identify challenges in this category, as well as some approaches towards solving them.

Usability and User Experience. This challenge is truly cross-cutting. The proposal would simply be infeasible if users are unable to specify workflows, provide rankings of various sites, etc. While we touch on this subject a little bit later in the paper, it is mostly orthogonal to our proposal, but, of course, essential.

In the rest of this paper, we will describe a challenge example for providing user-controllable workflows (Section 2) and then use it to describe some modeling and analysis challenges we as a community face before the Personal Web idea becomes a reality (Section 3). We also discuss how to begin solving these challenges (Section 4). We conclude in Section 5.

2 Motivating Example: Online Crib Shopping

We begin by proposing a (real-life!) challenge problem for Personal Web.

Consider the following scenario. Our (Canadian) user is six months pregnant and wants to purchase a baby crib. Quality cribs are durable but expensive, and take a while to get once ordered. So, she wants to try to buy a second-hand crib. The easiest way to get one is through a local online classified ads, such as craigslist.org, since she can go to the vendor in person and inspect it before making a decision. Her parents live in the US and frequently travel back and forth, so they can also look for local deals on used cribs, including their local craigslist.org, garage sales, etc. Our user also knows that quality cribs take 6 weeks to arrive when ordered, so she can only keep looking at used cribs for another 1.5 months. If that (soft) deadline passes, the user will have no choice but to go to a retailer that has cribs in stock and buy whatever they have – clearly not a good choice but might be the only option for meeting the hard deadline – having a crib once the baby arrives.

Here the goal is clear – to have a crib by a due date. But in addition, there is a set of preferences: (a) the user prefers a used crib but if none are available within 1.5 months, she will purchase a new one (although what if a perfect used crib becomes available within days of placing an order for a new crib. Can that order be canceled?); (b) the user wants to avoid shipping from the US in order to avoid customs delays as well as extra taxes. This means that if she buys the crib on a US website, she must remember to ship it to her parents' house (context-based preferences).

To accomplish this scenario, the user needs to interact with various services/sites:

- Research: product databases, review sites, user groups and forums.
- Purchase: auction sites, online classified ads, online retailers (and of course the related payment processing).
- Shipping: shipping estimator, shipping, truck rental.
- Utilities: currency converter, online spreadsheet, email, calendar, task lists.

She also needs to keep her parents up-to-date on the crib search, in order to avoid buying two cribs and coordinate travel dates in case her parents find a crib first.

To make this scenario into a Personal Web application requires, effectively, producing an orchestration for the above services, which allows prioritizing, context aware information, concurrency, and even undoing a finished task (such as attempting to cancel an order for a new crib if a used crib is found soon after). The orchestration should satisfy a number of properties, among them (a) a crib must arrive before the due date; and (2) at most one crib should arrive. For example, the last property is violated in the scenario where both the user and her parents independently find a good local deal on the same day and both decide to buy the crib on the spot to avoid losing the deal.

3 Challenges

As we mentioned earlier, we envision that many workflows can be "prebuilt" or community-shared – much like iPhone applets. However, there should be cases when a user may want to build their own customized and complex workflow, like in our example in Section 2. In this section, we describe specification and analysis challenges associated for providing support for automating complex user-created orchestrations.

1. The first obvious challenge is *specification* – specification of the desired outcome of the orchestration, available services, properties of the orchestrations, preferences, context, etc. The outcome of such specification should be sufficiently precise, so as to enable creation and reasoning about non-trivial service orchestrations, without the user knowing the technical details of service configuration or sequencing. And of course such specifications should be "average" user readable, without resorting to the use of formal logic.

2. Monitoring and statically analyzing correctness of orchestrations. Personal workflows are operationalized through orchestrations of user-level services. These orchestrations might be modified or recreated at runtime as services become unsuitable because of changes in user preferences and constraints (i.e., upon failing to secure the crib in time, the user may want to apply a completely different strategy such as borrowing a bassinet from a friend or deciding to do without a crib for the first few months of the baby's life).

1. Assuming it is possible to use the cloud or the user's own machine to check the workflow against the goals, how do we go about doing it without a major slowdown in workflow execution. What sorts of analyses are appropriate and how to make them scale? (the "compositionality" challenge)
2. Another set of questions for this challenge involves figuring out when a particular service can be substituted for another, whether it can be used in place of a collection of services in an orchestration, and suggesting which combinations are feasible. To enable personalized workflows, it is important to be able to aggregate services, resources, and content from multiple web services centering on the user, her tasks and context. In the crib purchasing workflow, the user may include several online retailers in her bookmarks in the search for a usable and inexpensive crib. Each of these retailers has their own specific shipping policies and methods. To achieve this simple workflow, several activities need to be performed: Different services and resources need to be identified. The compatibility of the interfaces of these services has to be verified. The services might need to be substituted dynamically when the user preferences or context change. For example, if the user realizes that her parents have bought a crib, the services in her workflow related to crib purchasing should be replaced with proper services enabling and coordinating crib shipping or the parents' travel schedule from US. Finally, the services must be composed periodically. (the "compatibility" challenge)

3. Repairing orchestrations. Errors occur in most software applications, but they are unavoidable in web-based applications. They can happen because some partner went down, Internet became unavailable or the logic got violated. In our example, the user wants to buy exactly one crib. However, it could happen that both the user and her parents buy a crib on the same day, violating this user requirement. Repairing user-created faulty orchestrations is clearly needed.

Of course, the users should be able to specify orchestrations and their properties, and once an orchestration is deployed, monitor and repair the orchestration to make sure that their goals are accomplished. We also envision some form of an "orchestration dashboard", where users can create new orchestrations and check state of orchestration instances.

4 Approaches

In this section, we discuss some ideas to approach the challenges identified in Section 3.

1. The Specification challenge. One way to specify workflows is to avoid such specification altogether! Specifically, we may want to synthesize complex workflows automatically, based on the expression of user intent. Effectively, intent is a declarative specification which is then turned, by synthesis, into an operationalized workflow. One way to approach it is to adapt existing work on configuring personal software using goal models [5,6].

In this work, an i* goal model [11] is used to specify a set of possible user goals, how they interact and how they decompose into simpler goals. Each low level service configuration setting contributes, to different degrees, to different higher level goals. For example, in our scenario, the user has a goal ”Minimize cost”. In this case, each product provider service can specify how much its configuration settings support this goal (e.g., selecting ”used” rather than ”new” items will contribute more to this goal). Thus, by allowing the users to identify their goals, the optimal service configuration settings can be determined automatically. Furthermore, this can be used to identify when user goals are unsatisfiable (since no settings exist to satisfy them all) and hence require the user to modify their goals.

The goal model also supports AND/OR decomposition and this can be used to automatically define a service sequence. An AND decomposition step defines the set of subgoals that must be satisfied to satisfy a goal. This can be used to constrain the possible service invocations required and the input/output dependencies between these services constrain the possible sequencing of them. An OR decomposition step specifies a set of alternative subgoals and this leads to an interaction with the user to refine their intent by choosing an alternative. For example, the goal ”Deliver product quickly” can have alternatives ”Deliver within 2 weeks”, ”Deliver within 6 weeks”.

2. The Monitoring and Analysis challenge. To address the compositionality challenge, we begin by looking at approaches for verifying web service orchestrations. Such verification can be done statically or dynamically, e.g., [1–3, 9] and can be relatively easily adapted to user-created orchestrations. However, the distinction between the two approaches is that the user-created ones are much more agile – once anything in the orchestration changes, the whole analysis may have to be redone, and the user is forced to be aware of it as the performance of her system deteriorates. For example, the verification may concern basic functional properties such as “the crib is eventually bought” or the more complex ones such as “the crib is usable and has a reasonable price”, “only one crib is bought”, or “the crib has been delivered within an acceptable time period”. However, periodic application of these techniques in service compositions that evolve over time, where services are frequently added, removed, or revised, is unrealistic. To ensure correctness, we need to design service compositions in a way that verification results can be reused across evolutions, i.e., to allow *regression verification* [4].

To achieve this goal, we propose to exploit composition design patterns [8] and orchestration algorithms that make verification *change-aware*.

To address the substitutability challenge, we aim to create support for compatibility and substitutability of web services. We propose to do this by investigating similarity measures between behavioural models [7] to identify candidate services to replace a service in use when it becomes unavailable or unsuitable due to evolving needs or a change in the context. Further, we require composition techniques for combining these services. This is a non-trivial activity as it requires integrating both behaviour and data assets of services while preserving their desirable properties and ensuring that interactions between services within a composition do not lead to unforeseen and highly undesirable side-effects.

3. Repairing orchestrations. The goal of this activity is to use semantic information about services involved to try to fix the problem discovered for an orchestration. For example, consider the case when two cribs were purchased at the same time. The problem can be fixed if the user sells one of the cribs, as this returns the orchestration to a state where the user has just one crib. The user may have preferences as to how to accomplish this, e.g., she may choose to sell the most expensive crib, or maybe the one that will arrive last, or the one bought by her parents.

In another scenario, imagine that the parents bought a second-hand crib, but decided to ship it instead of transporting it themselves. Due to a backlog at the customs office, the crib will not arrive by the required date (violating the "the crib has been delivered within an acceptable time period" property). In this case, we can repair the orchestration by suggesting that the user buy a new crib locally (which leads to the satisfaction of the violated property), while selling the other crib when it arrives (to avoid violating the "at most one crib" property).

In the first scenario, we suggested actions that compensate executed actions, leaving the orchestration in a state that does not violate any user requirements. In the second scenario, we also suggested the execution of new activities that lead to the satisfaction of user requirements. We have explored the idea of property-guided recovery in the context of traditional web applications [9, 10], where both the orchestration and its properties are defined by the application developer, but recovery plans are computed for individual execution traces. This framework can be adapted to the Personal Web paradigm. However, as we mentioned earlier, the success of such an approach depends on the ability of end users to specify correctness properties. Also, in the approach of [9, 10], activity compensation and cost are statically defined. In order to move our approach to the Personal Web, compensation and its cost should be user-specified (e.g., to account for cases where some users pay smaller fees for a transaction cancellation, be that for a stop payment or for cancelling a flight). This is theoretically possible, of course, but so far we are not aware of technology that would allow us such dynamic, user-centered compensation definition and configuration.

5 Summary

In summary, the dream of Personal Web seems achievable, provided a number of challenges are met. Many of these are technological, and we have no doubt about their success. Some others need advanced techniques and thus require research. We tried to argue that problems of modeling and analysis of end-user created orchestrations are important, and solutions to them are possible. However, success of the whole endeavor depends on several computing fields and, most crucially, on creating a good user experience, especially in support for specifying desired orchestrations and their properties.

We look forward to collaborating with others on various aspects of problems brought forth by Personal Web.

References

1. H. Foster, S. Uchitel, J. Magee, and J. Kramer. "Model-based Verification of Web Service Compositions". In *ASE '03: Proceedings of 18th IEEE International Conference on Automated Software Engineering*, pages 152–163. IEEE Computer Society, 2003.
2. X. Fu, T. Bultan, and J. Su. "Analysis of Interacting BPEL Web Services". In *WWW '04: Proceedings of the 13th International Conference on World Wide Web*, pages 621–630, May 2004.
3. S. Hallé and R. Villemaire. "Runtime Monitoring of Message-Based Workflows with Data". In *ECOC '08: Proceedings of the 12th IEEE Enterprise Distributed Object Computing Conference*, pages 63–72, 2008.
4. T. A. Henzinger, R. Jhala, R. Majumdar, , and M. Sanvido. "Extreme Model Checking". In *Verification: Theory and Practice*, volume 2772 of *Lecture Notes in Computer Science*, pages 332–358, 2004.
5. S. Liaskos, A. Lapouchnian, Y. Wang, Y. Yu, and S. Easterbrook. "Configuring Common Personal Software: a Requirements-Driven Approach". In *RE '05: Proceedings of International Conference on Requirements Engineering*, pages 9–18, 2005.
6. S. Liaskos, A. Lapouchnian, Y. Yu, E. Yu, and J. Mylopoulos. "On Goal-based Variability Acquisition and Analysis". In *RE '06: Proceedings of International Conference on Requirements Engineering*, pages 76–85, 2006.
7. S. Nejati, M. Sabetzadeh, M. Chechik, S. Easterbrook, and P. Zave. "Matching and Merging of Statechart Specifications". In *ICSE '07: Proceedings of the 29th International Conference on Software Engineering*, pages 54–64, 2007.
8. S. Nejati, M. Sabetzadeh, M. Chechik, S. Uchitel, and P. Zave. "Towards Compositional Synthesis of Evolving Systems". In *FSE '08: Proceedings of SIGSOFT International Conference on Foundations of Software Engineering*, pages 285–296, 2008.
9. J. Simmonds, S. Ben-David, and M. Chechik. "Guided Recovery for Web Service Applications". In *FSE '10: Proceedings of SIGSOFT International Conference on Foundations of Software Engineering*, pages 1–10, 2010. to appear.
10. J. Simmonds, S. Ben-David, and M. Chechik. "Monitoring and Recovery of Web Service Applications". In *Smart Internet*, pages 1–40. Springer, 2010. to appear.
11. Eric S. K. Yu. "Towards Modeling and Reasoning Support for Early-Phase Requirements Engineering". In *RE '97: Proceedings of IEEE International Symposium on Requirements Engineering*, pages 226–235, 1997.

The Personal Web: Personal Observations & the Challenge for Financial Services

Ian Graham

Bank of Montreal

Managing Context

Notification Design Challenges for the Personal Web

Jin Li and Jimmy Lo

IBM Canada Ltd., Toronto, Ontario, Canada

jinli@ca.ibm.com, jimmylo@ca.ibm.com

Abstract

Based on research conducted to improve the design and management of notifications in the Jazz collaborative development environment and a CASCON workshop on notifications, we describe a set of notification design challenges for the Personal Web.

1 Introduction

There has been recent Human-Computer Interaction and related research on the topic of managing interruptions [7]. Interrupting users with many non-critical events, for example, has been shown to lower productivity and cause stress and frustration [8]. In this paper we use the more neutral term "notification" to mean the same thing as a system delivered "interruption." The two main components to interruption management that have been studied, albeit largely independently, are: (1) timing of interruption delivery based on user context [4]; and (2) presentation format of the interruption based on user context [1, 10]. In a pure timing-based model, notifications are queued up until it is a "good time" to interrupt the user, and all of the queued notifications are delivered at once. In a pure presentation-format based model, notifications are always delivered right away, but the way that the notification is presented, in particular the degree of intrusiveness is adjusted based on context.

A user's context generally refers to what the user is doing at a given moment, and at its highest level reflects whether the user is busy or not. The two primary ways to determine context are: (1) Automatically (or system determined), where the system attempts to detect when users are interruptible or not [6, 12], and (2) Manually (or user determined), where users declare their preferences as to when they can be interrupted, and the system interrupts users accordingly. Context is more involved than simply the degree to which the user is busy doing an activity. Knowing when to interrupt should also depend on the importance of the interruption content, which is referred to as the interruption's relevance or utility.

The Personal Web is the next instantiation of the *Smart Internet Research Initiative* [9]. The Personal Web focuses on the person *ME* as the center of gravity of Web integration. The objective of the *Personal Web* is to empower *ME,* as a common Internet user of generally limited technical skills, the autonomy and ease of control in assembling and aggregating integrate-able elements across the Web for a particular *sphere of context* of my concern. Making it simple for users to visualize and interact with the results of the user-sovereign Web integration would make ME productive and delight ME. Notification design optimized for the sphere of context of my concern is pivotal in enabling a simple user model, making ME productive by directing my attention to the most important and relevant information. For the Personal Web, it is important that notification delivery is adapted to factors such as device and location. In the Personal Web, it is not just about user's current task but also the current sphere of context, which the user has manually defined and the system has ongoing implicit knowledge of.

This is ideal for providing information about the user's context.

The main contribution of our work is a set of notification design challenges for the Personal Web. In Section 2 that follows, we will preliminary results from our recent research and workshop on notifications. In Section 3, we will describe the notification design challenges for the Personal Web. We will conclude this paper with proposed next steps.

2 Research Conducted

We have recently done an effective notification design research on the Jazz collaborative development environment (CDE) [5]. The high-level goal of this research is to investigate better collaboration support through appropriate awareness of team members' activities in the Jazz Web UI environment. The research specifically aims to improve collaboration through interrupting users at context-appropriate times with relevant information in context-appropriate presentation formats.

We have also hosted a CASCON 2009 workshop on notifications, which was to explore the effective design of notifications in various collaborative / social environments, including but not limited to collaborative development environments and social networking sites.

In this section, we will briefly summarize the research conducted and their preliminary results.

2.1 Notifications in Jazz CDE

Jazz is an IBM initiative to help make software delivery teams more effective. It is a CDE designed to transform software delivery, making it more collaborative, productive and transparent for development teams using the agile development process. Instead of an integrated development environment (IDE) such as Eclipse which focuses on supporting an individual software developer, a CDE puts the team first with the assumption that team productivity will increase [2].

Agile development uses a less structured development process relative to more traditional software development, and thus has a greater need for awareness of team member activities. Agile encourages rapid peer-to-peer communication, relying much less on official top down communication channels, such as team meetings. Jazz supports informal communication and light-weight awareness through a number of different mechanisms, which include notifications of events that are fired based on many user activities (such as completing a work item) as well as system activities (such as the outcome of a build).

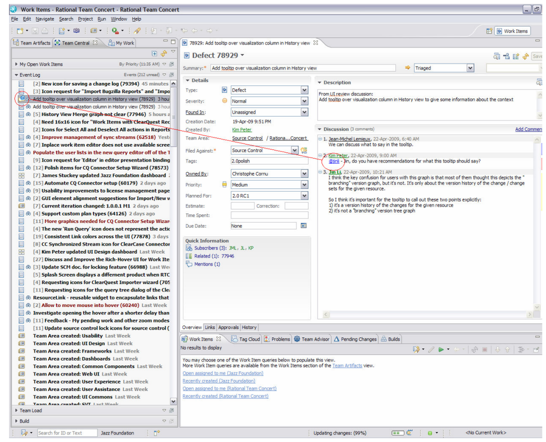

Figure 1: The Jazz CDE, showing the events list (left), linking to work item data (right).

We took a scenario based design approach [3] and conducted a focus group user study with eight representative Jazz users to validate the scenarios and generate design requirements. We conducted a three-session focus group study, each session lasting 1.5 hours, with eight representative participants, all IBMers who use Jazz.

Our high-level goal was to elicit requirements. The specific goal of the first session was to present and validate the As-Is scenario (which reflected our assumptions about the current problems with notifications in Jazz) and to document specific pain points experienced by our participants. After the first session we modified the To-Be scenario (which addressed those problems with specific design solutions) reflecting what we had learned. The goal of the second session was to validate the revised To-Be, and to prioritize the pain points elicited in the first session. From these two sessions we learned that our scenarios were valid, although some minor adjustments were required. Our participants told us that they are over exposed to notifications in the Jazz CDE, resulting in poor awareness of team activities. The pain points that emerged in the first session and that were ranked as most problematic in the second session were as follows:

1. Awareness of changes by others that affect my current work
2. Prioritizing all incoming communication related to my current work

3. Notifying others who are (may be) affected by changes in my work
4. Blocked communication path to others

Our main findings can be categorized into the following four themes:

A) Individual needs come first, team needs second.
B) Notification customization UI design presents unique opportunities and challenges.
C) Notification UI should not block users: control must remain with the user.
D) Don't want technology "watching you."

We then conducted brainstorming sessions within our research team to prioritize the requirements and explore design ideas, generated design sketches, and solicited input in the third focus group session. The key design goal that emerged was to create a central place to view and prioritize work items and high-priority events. Finally, we created a proof of concept working prototype for the Web version of Jazz CDE. The prototype supports many of the features and design elements described above.

2.2 CASCON Workshop

In CASCON 2009, we hosted a workshop on notifications. We spent the first 15 minutes for participant self-introduction, and we had 30 workshop attendees. Two 20-minute short panel presentations followed: we presented our Jazz CDE research; a product manager at Google presented their work on notifications for mobile applications. At the end of the presentations, there was a 30-minute panel questions and discussion. After a 15-minute break, we had four break-out groups and each group had 45 minutes to discuss and answer a broad range of questions and issues, including the following:

- What is the design space for the presentation styles of notifications?
- How should notification content (the message) be best matched to a presentation style?
- To what extent does the effectiveness of notifications (presentation styles and notification content) depend on the particular kind of environment or application?
- What is the state of the art in commercial environments and applications?
- What is the state of knowledge on notifications from the research literature?

Finally, we regrouped and had 30 minutes to wrap up the workshop, with each group presented a summary of their discussion.

We intend to write up a separate full workshop report, but here is a summary of the key findings:

- Participants told us that they are indeed being interrupted by too many forms of notifications, from email, RSS feeds and Facebook status updates to alarm clocks, microwave buzz and fire alarms.
- Participants concurred that good notification design would be a critical component for success in the Web and mobile applications space.
- Participants were divided on intrusiveness of notification presentation styles. Some prefer more intrusive presentation style, such as a pop up dialog box, while others prefer non-intrusive presentation style – regardless of notification content (the message).
- Participants were also divided on whether notifications should be viewed and controlled from a central place, e.g. a dedicated notification bar.
- Participants voiced concerns about looking at notification design purely from a software only perspective. A systems approach might be more appropriate, where notifications are delivered within the user's broader environment. For example, instead of a "You've got mail" notification showing in one's email program, the user's mobile phone vibrates; instead of a buzz via Instant Messaging (IM) chat window, the user's phone rings, interrupts the user, and when the user picks up the phone it tells the user that there's an urgent IM chat request from the spouse.
- Participants believed that it is important to design the system so that it knows when to override user's notification preference, for safety and other critical reasons. For example, even when users have set to be "Do not disturb" state with respect to notifications, the system must interrupt and alert users to a "fire alarm" equivalent notification.

3 Design Challenges

Much of current interruption management research has focused on how to reduce notifications (thus user interruptions), in the hope that it would increase user's productivity. Since the Personal Web focuses on the person *ME* and is to empower *ME* and has implicit knowledge of the user's current sphere of context of my concern, we propose that we turn the research question around and ask how we could increase the useful notifications that would make ME smarter.

From our research to date, most interesting and unique design challenges for the Personal Web are the following:

#1 Individual needs come first, team needs second. Despite the intention behind a Jazz CDE to support and benefit a team, the team is not the first priority for users. Design should focus on benefits to the individual, with the expectation that improved notification and awareness for the individual will in turn lead to a productivity improvement for the team. Although the Personal Web is also a part of the larger social Web and ME being the center of the Personal Web, the design challenge is that notifications are only there to serve ME, not my social network or making ME look good or social among my friends. Research is also needed to investigate how users could use notifications as a navigation mechanism in the Personal Web.

#2 Don't want technology "watching you."
Jazz CDE users don't want to feel like they are being watched. Notifications can serve different purposes for different roles. Some developers choose not to follow team process (e.g., deliberately do not change the status of work items to "in progress") because this often generates many notifications that only help others to track their progress. For the Personal Web, the design challenge is not just about user privacy (which may or may not be a major concern in the future), but also complete control, transparency and traceability of my notifications.

#3 A systems approach is more appropriate.
CASCON workshop participants voiced concerns about looking at notification design purely from a software only perspective. A systems approach might be more appropriate, where notifications are delivered within the user's broader environment. For the Personal Web, the design challenge is that notification delivery should be adapted to factors such as device and location -- the user's current sphere of context of my concern and the user's current environment.

#4 To centralize notifications or not
CASCON workshop participants were also divided on whether notifications should be viewed and controlled from a central place, e.g. a dedicated notification bar. For the Personal Web, the design challenge is whether a centralize notification view, or a distributed notification UI mechanism where it brings notifications closer to the user's sphere of context of my concern and the user's current environment works better.

4 Conclusions

Based on our recent research and preliminary findings, we have described a set of unique design challenges for the Personal Web. The Personal Web sharpens the focus on the broader interruption management research to a few unique issues. We believe good notification design would be a critical component for the success of the Personal Web. We will start investigate solutions to some of the design challenges.

Acknowledgements

The authors would like to thank IBM CAS Canada for its support in this research; thank members of the Smart Interaction POC team for their collaboration and notification design prototyping.

About the Authors

Jin Li is a user experience lead at IBM Canada Ltd. His work encompasses the full software design lifecycle. In particular, he gathers, analyzes, and translates user requirements and usage scenarios into software and user interaction designs for application development tools and Web. Currently, he is responsible for the user experience of Rational Rhapsody and systems engineering tools. Jin holds a Masters of Science in Computer Science, with a Human-Computer Interaction option, from the University of Toronto. He is a member of ACM and UPA. Jin can be reached at jinli@ca.ibm.com

Jimmy Lo is a Research Staff Member at IBM Canada Centre for Advanced Studies Research. Currently, he is responsible for managing fellowship projects with Smart Interaction technology theme. Jimmy is a Professional Engineer and holds a Bachelor of Engineering Science from the University of Western Ontario. Jimmy can be reached at jimmylo@ca.ibm.com

References

[1] Avrahami, D., Hudson, S.E.. QnA: Augmenting an instant messaging client to balance user responsiveness and performance. In: Proceedings of ACM CSCW 2004, 515-518 (2004)

[2] Booch, G., and Brown, A.W. (2002). Collaborative Development Environments. Online: http://www.booch.com/architecture/blog/artifacts/CDE.pdf, retrieved 10/15/2009.

[3] Carroll, J.M. Ed. (1995). Scenario-based design. John Wiley. New York, New York.

[4] Chen, D., Vertegaal, R.: Using mental load for managing interruptions in physiologically attentive user interfaces. In: Extended Abstracts ACM CHI 2004, 1513-1516 (2004)

[5] Chignell, M., Cordy, J., Ng, J., Yesha, Y. (Eds.), The Smart Internet, Chapter 3, Springer (2010).

[6] Horvitz, E., Apacible, J.: Learning and reasoning about interruption. In: Proceedings of ACM ICMI'03, 20-27 (2003)

[7] Interruptions in Human-Computer Interaction, http://interruptions.net/

[8] Mark. G., Gudith, D., Klocke U.: The cost of interrupted work: more speed and stress. In: Proceedings of ACM CHI 2008, 107-110 (2008)

[9] Ng, J.: The Personal Web. In Proceedings of Personal Web Symposium, CASCON 2010.

[10] Oberg, B., and Notkin, D.: Error reporting with graduated color. In: IEEE Software, 9(6):33–38. (1992)

[11] Sample, B.: A Software Design Approach. In *Proc. Of the 15th Software Conference*, pages 40-50 (1995).

[12] Van Dantzich, M., Robbins, D., Horvitz, E., Czerwinski, M.: Scope: Providing awareness of multiple notifications at a glance. Proceedings of AVI (2002)

Managing Dynamic Context to Enable User-Driven Web Integration in the Personal Web

Norha M. Villegas, Hausi A. Müller
University of Victoria, Victoria, British Columbia, Canada

Joanna Ng, Alex Lau
IBM Canada Lab, Toronto, Ontario, Canada

October 15, 2010

Abstract

The *Personal Web* is the people-centric instantiation of the Smart Internet where information systems, services and web content are articulated by users according to their matters of concern. To realize the vision of the Personal Web, the Smart Internet requires infrastructure to support the user in the integration of personal data and the composition of personal services within a highly dynamic context that constitutes the user's *Personal Web Sphere*. To address these requirements, we propose a user-driven context management framework, built on the top of the basic enabling infrastructure of the Personal Web, to support users in the run-time modification of *personal context models*. The core of our proposal is the management of monitoring concerns by implementing feedback loops, where the user acts as the planner of the controller to adapt the monitoring strategy by means of using web interactions to modify the personal context models. These context models, deployed at three different levels of abstraction, represent monitoring concerns by defining abstract types of contextual entities, the relationships among them and the interactions that the user can instantiate to drive web integration.

1 Introduction

The *Personal Web* is a concrete realization of the Smart Internet that focuses on the user as the center of web integration [7]. Thus, Smart Interactions are required to support, from a people-centric perspective, the discovery, aggregation and delivery of resources from the Internet. Moreover, Smart Services must provide the infrastructure required by these interactions to assist users in web integration [8]. The vision of the Personal Web is to enable regular web users (i.e., people with non-specific technical skills) to control the integration of web resources according to their personal matters of concern. The semantics that define relevant types of web resources, the relationships among them, and the way how the user interacts with these entities is defined as the *Personal Web Sphere*. In this way, the user's personal web sphere defines the environmental context that affects the web integration performed by the user.

Due to the complex, volatile and transient nature of context, the Personal Web requires innovative approaches to assist the user in managing the information that can affect the interactions with web resources. To address this complexity, we exploit feedback loops and context-awareness techniques to propose a context management framework that supports and empowers the user in integrating web resources (web context entities) according to personal matters of concern. Our approach, built on the top of the Personal Web's enabling infrastructure proposed by Ng [7], supports the user in the modification of personal context models to drive the integration of web resources at runtime. These personal context models, modified by the user through simple but meaningful web interactions, feed our context management infrastructure back to discover personal web-enabled context instances that can be integrated to the user's personal context models, feeding back the context management system.

In this paper, we propose a user-driven context management framework to discover personal context entities and enable the user in the dynamic integration of these entities to her personal web sphere. With our approach, we support four of the five conceptual principles and the whole set of technical principles defined by the vision of the Personal Web [7]. With respect to the conceptual principles, (i) the user is the center of the web integration acting as the planner of our context management controller; (ii) the scope and semantics of the web integration is controlled by our context management framework according to personal context models (the user's personal web sphere); (iii) context-aware interactions between the user and entities of the real world are controlled by the user to accommodate context management concerns according to every day situations, by means of adapting personal context models; and (iv) these context management concerns are persistent in the form of personal context models that are re-used every time the user interacts with the corresponding personal web sphere. Regarding the technical principles, (i) having personal web-enabled context entities (web resources), the interactions defined by personal context models enable the user to compose desired entities into her personal web sphere readily; (ii) as our framework is based on the basic enabling infrastructure of the Personal Web, characteristics of design for open integration and simplicity of operations for web integration are preserved; and (iii) our modeling approach, deployed on the Personal Web's three layered meta-model, supports the instantiation of interactions between the user and instances of context entities, as well as the relationships among these entities that are relevant for the user's interactions.

The remainder of this position paper is organized as follows. Section 2 describes the scenario used throughout the paper to explain the components of our framework. This scenario will drive our validation. Section 3 explains our context management framework, as well as the integration of the basic enabling infrastructure of the Personal Web and its three-layered meta-model into our proposal. Finally, Section 4 discusses important aspects of enabling user-driven web integration by means of dynamic context management.

2 A Motivating Scenario

Imagine a user who is working on her web-based calendar. Her current concerns are related to two particular events, the anniversary dinner with her husband and her next business trip. Her *personal web* is required to assist the user in different context-aware, on-line activities concerning each event. Whenever she works on a particular event, the *smart Internet* supported by our *dynamic context management infrastructure* ought to suggest personal web-enabled data and services according to her matters of concern. These personal web-enabled elements, which we call *personal context entities*, are displayed on her web either because they are already within her personal web or because they are related to her current situation. Moreover, personal context entities previously existing in her personal web have a defined relationship with the user. On the contrary, personal entities suggested according to her matters of concern only will become part of her personal web when the user defines a relationship with them. These relationships are defined by the interactions between the user and the personal web enabled context entities discovered by our context management infrastructure at run-time.

Imagine the user is working on the anniversary dinner event. Her personal web displays linked data such as the invitation email sent by her husband to her. This email is a personal context entity (already existing in her personal web) that contains context information such as the restaurant where the dinner will take place. Based on her *personal web context models* enabled by our context management infrastructure, it is possible to suggest web services such as on-line shopping services according to the user's preferences and situation. Suppose she interacts with her favorite designer's fashion boutique web service. As the woman is already registered in the boutique's system, this web service previously linked to her personal web displays a custom dress catalog taking into account the current user's matter of concern (e.g., context information related to the dinner, location and clothing preferences). Then, she places an on-line order buying that new fancy dress she wants to wear for the anniversary dinner. As she pays using her credit card, the on-line banking service is also added to her personal web for future on-line shopping activities. Finally, while the user browses the boutique's web site, she interacts with some of the boutique's partner web services to order shoes and accessories suitable for her nice new dress. These URL resources are also added to her personal web as new *personal linked services*.

Because users' concerns vary according to current situations, one user can have more than one personal web instance. Suppose our user gets a calendar alert concerning her next business trip. Using her personal web models for traveling and shopping concerns, the smart Internet infrastructure suggests personal web services to buy air tickets, book a hotel and rent a car for this trip. As the woman browses the web, new personal context entities are added to her personal shopping web. Nevertheless, existing personal linked services such as her credit card payment service must be integrated to the current shopping activity. It is important to point out that shopping concerns differ for each event. Thus, a personal web infrastructure supported by the smart Internet is required to assist the user in her web experience taking into account her current situation.

In this motivating scenario the user acts as the main controller in the definition of her personal web. Supported by our context management infrastructure built as part of the smart Internet infrastructure, the user will be able to define new relationships with personal web-enabled context entities at run-time. This motivating scenario is used throughout the following sections to explain how, exploiting feedback loops and context awareness techniques, we can manage dynamic context to assist users in driving web integration as proposed by the Personal Web.

3 Dynamic Context Management for the Personal Web

Figure 1 illustrates the dynamic context management framework we propose to assist users in driving web integration in the Personal Web. Our approach is based on the feature-oriented reference framework we proposed for evaluating and implementing context management infrastructure, as required by the Smart Internet, and on the context management infrastructure we proposed for realizing dynamic monitoring in SOA governance [12, 11]. As we discussed in [12], the first big challenge related to context-awareness in the Smart Internet is the *identification of relevant context* for a specific set of interactions or user's matters of concern. The second challenge is then the *modeling* of these requirements in such a way that the representation of context is able to adapt itself dynamically as the application's environment and the user's concerns change. Once relevant context and context management objectives are modeled, the third challenge is to *design and implement* a dynamic context management infrastructure able to support the gathering, handling, and exploitation of context according to the model.

To propose this framework, we hypothesize that web integration can be addressed by integrating the user into a context management feedback loop. In this feedback loop, the user acts as part of the controller, specifically as the planner, to modify models of the real world that support the run-time instantiation of existing or new types of context entities within the user's Personal Web Sphere. To realize context modeling, we integrate the context taxonomy, the feature-based models and the context meta-model we proposed in previous works [11, 12] with the three layered meta-model envisioned as part of the basic infrastructure of the Personal Web [7]. To realize dynamic context management, our user-centric feedback loop supports the discovery and monitoring of relevant entities, and the adaptation of personal context models according to the user's interactions with these entities.

3.1 Modeling Personal Context Entities

Context modeling is an important component of the context information life cycle in the Personal Web. Personal context models represent the relevant aspects of the interactions between the user and entities of the real world, as well as the relationships among these entities that can affect the user's interactions. In our approach, entities of the real world that are relevant for the user's personal web sphere are known as *personal context entities*. To represent personal context entities, we define context models on three different levels of abstraction based on the three-layered meta-model proposed for the Personal Web [7]. These three layers correspond to the concept layer, the model layer and the instance layer. According to Fig. 1, *the concept layer* is composed of a foundational feature-based meta-model that defines the types of context entities and the relationships among them according to the general context taxonomy we proposed for the Smart Internet [12]. This meta-model is complemented with domain specific ontologies to instantiate the domain specific personal context models that define *the model layer*. Finally, the actual instances of the personal context models, which are derived from the interaction between the user and personal web enabled context entities (web application layer), constitute the *instance layer*.

3.1.1 The Concept Layer: The Foundational Context Representation

The feature-based meta-model for dynamic context management and the context taxonomy we proposed in [11, 12] provide the abstractions for the instantiation of personal context models. This feature-based meta-model defines the abstract context data types and the relationship among them. Relationships among context entities support the instantiation of both the relationships between the user

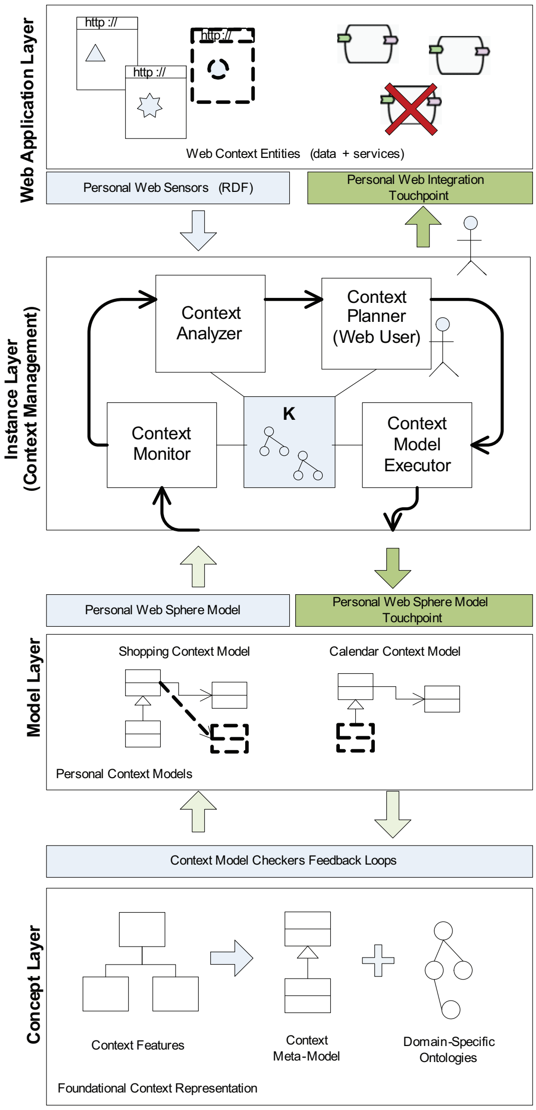

Figure 1: User-driven dynamic context monitoring framework for the Personal Web. The user acts as the planner of the feedback loop by modifying context models that drive the discovery and monitoring of personal data and personal web services according to the user's concerns. Personal context models are defined at three levels of abstraction as proposed by the Personal Web.

and entities of the real life, and the relationships among these entities. This meta-model also defines the abstractions that represent relevant properties of these entities. The abstract definition of the relationships between the user and context entities will enable the linking of personal data and personal services at the instance level. To realize the instantiation of personal web context models from this abstract representation, domain-specific ontologies complement the concept layer to express semantic dependencies between instances of context information in particular web domains [10, 3]. For our motivating scenario, the concept layer defines the abstract types required for the instantiation of the personal context models that represent the context entities relevant for the personal shopping web sphere and the personal calendar web sphere of the user. The meta-model together with domain-specific ontologies also defines the instantiation of valid interactions between the user and the relevant context entities. Instances of these interactions are a "debit" interaction with the "credit card" service entity, a "buy" interaction with a "dress" product entity through a "shopping car" service entity, and an "attend" interaction with an activity context entity such as the anniversary dinner or the business trip.

3.1.2 The Model Layer: Personal Context Models

From the abstract representation of context entities and the relation among them supported by web domain-specific ontologies in the concept layer, the model layer instantiates basic domain-specific personal context models required by the user's personal web sphere. These models are used to represent actual types of personal context entities that can be relevant for the user, the valid interactions between the user and these entities, and the relationships among these context entities required to support the user's interactions. Moreover, these concrete models define the information that characterizes the situation of entities the user interacts with. Personal context instances that define the user's personal web sphere in the instance layer are derived from personal context models of this layer. Revisiting our motivating scenario, a particular instance of personal context model for a shopping concern should include context entities such as "product", "on-line vendor", "financial institution", "credit card", "on-line shopping car", "calendar event"; user interactions such as "buy" a product, "attend" an event, "debit" from a credit card; and relationships among types of context entities such as "partnership" among different on-line vendors. Supported by our monitoring feedback loops, the user drives the adaptation of these personal context models according to her web integration concerns.

3.1.3 The Instance Layer: Personal Context Entities

As depicted in Fig. 1, the instance layer is composed of the context management feedback loop including the knowledge base. Personal context instances derived from the model layer are part of the knowledge base of the user's personal web sphere. Personal context instances represent entities of the real world and the way how the user interacts with them. They also model the relationships among these entities that are relevant for the user's interactions. Examples of context instances from our motivating scenario are the "dress" that the user is buying, the "fashion store" where she is placing the purchase order, her "financial institution", her "credit card", the "shopping car" where the dress, the shoes and the accessories are registered in, the "anniversary dinner" and the "business trip". Relationships and interactions are instantiated in the same way as proposed in the model layer but between concrete context entities and the user.

As our framework is based on the enabling infrastructure of the Personal Web [7], the identification personal web-enabled context instances, the relationships among them, and user's interactions are based on the Resource Description Framework (RDF) [1, 5]. Instances of actual implementations of RDF are the OpenLike protocol supported by Google and the Open Graph Protocol (OGP) proposed by Facebook. Both protocols focus on capturing relationships among users by defining "like" interactions with web entities [4, 13]. Context entities are represented as web object nodes in an RDF graph (e.g., data entities, services). In the same way, the edge between a pair of entity nodes can represent a semantic relationship between these two context entities. These relationships are driven by the user through personal interactions. New interactions link new data or new services to the user's personal web sphere. The semantics of these personal context instances are defined using RDF properties. Moreover, mechanisms based on RDF can be used by personal web sensors to discover relevant context entities according to the context definition of the user's personal web sphere.

3.2 User-Driven Monitoring Feedback Loops

By dynamically upgrading personal context models (cf. the model layer) that support the instantiation of context entities and user's interactions (cf. the instance layer), the user controls her personal web sphere to integrate new personal data and personal services. Whenever the user interacts with entities within a particular web domain, concrete context entities and interactions instantiate the personal context model to set up the user's personal web sphere accordingly.

Inspired by the autonomic element of autonomic computing and its application to the engineering of adaptive systems [2, 6], we propose a context management infrastructure where the user acts as the central part of the controller. By selecting interactions from the available set of interactions with entities of the real world, the user acts as the *planner* of the feedback loop. Interactions performed by the user are contextual events that triggers the modification of personal context models. These modifications are performed by the *context model executor* according to the user's interactions. When interactions imply the addition of new types of context entities to the personal model (cf. model layer), the *context model checkers* control the instantiation of the general meta-model (cf. concept layer). While the user is browsing the web, the *context monitor* uses the corresponding personal context model to discover, through "RDF sensors", potential personal web-enabled context entities according to her matters of concern (e.g., the anniversary dinner and the business trip). Then, the *analyzer* correlates information from context entities with properties of the corresponding personal context model to identify a set of relevant interactions between the user and these entities of the real world. The user can also act as part of the analyzer by selecting as relevant, interactions automatically identified as non-relevant.

4 Discussion and Future Work

Our ongoing research focuses on the validation of available technologies for the implementation of our context management framework. Empirical evaluation is also required to assess the impact of the proposed approach in supporting the vision of the Personal Web.

This paper proposes a dynamic context management approach to assist the user in the integration of personal data and personal services into her personal web sphere. This way, we contribute to the realization of a *dynamic Personal Web Link* by modifying, at run-time, the models that drive the discovery of personal web-enabled entities (personal context entities) and guide the interactions between the user and these context entities. With respect to modeling aspects of personal web domains, modeling artifacts are implemented at three different levels of abstraction as proposed for the Personal Web [7]. Feedback loops play an important role to guarantee consistency among these layers of abstractions. Between the two lowest layers, the model and the concept layer, context model checkers (implemented as feedback loops) monitor transformations in personal context models to check the properties of the resulting models with respect to the conceptual properties represented by general context-meta models and the corresponding domain-specific ontologies. Between the instance layer and the model layer, the relation is twofold. First, a top-down interaction enables the context model executor from the instance layer to trigger transformations of the personal context models defined by the model layer. Second, changes in personal context models modify the set of potential personal web-enabled entities to be monitored by the context monitor. Whenever a personal context model transformation is performed, context model checkers perform the validation against conceptual models defined by the bottom layer. It is important to point out that not all the interactions performed by the user will trigger transformations in personal context models. The user could integrate new personal data or services that are instances of existing types in the actual personal context models.

With respect to architectural concerns, every element of the framework should be designed independently to support distribution. Enabling infrastructure of the Personal Web is required to be deployed either at the server domain, at the client domain or distributed on different processing nodes. For supporting architectural decisions, control-based reference models and reference architectures provide valuable guidelines for the implementation of dynamic monitoring infrastructure [11].

Finally, more research on self-adaptive service-oriented systems is required to leverage the user experience in the Personal Web. For instance, dynamic sensing infrastructure for discovering smart services based on dynamic service selection must be integrated to provide the user with a richer set of resources from the Smart Internet, beyond personal web-enabled context entities discovered from web pages [11, 9].

Acknowledgements

This work grew out of collaboration with colleagues at IBM Toronto Center for Advanced Studies, CA Labs, and Carnegie Mellon Software Engineering Institute. We are deeply indebted to many of our friends and students, who contributed significantly to our appreciation for the personal web and understanding of context-awareness. In particular, we would like to thank Gabriel Tamura, Grace Lewis, Sowmya Balasubramanian, Ron Desmarais, Priyanka Gupta, Ishita Jain, Stephan Jou, Przemek Lach, Atousa Pahlevan, Dennis Smith, Qian Yang, and Qin Zhu for their continued collaboration and support. This work was funded in part by the National Sciences and Engineering Research Council (NSERC) of Canada (CRDPJ 320529-04 and CRDPJ 356154-07), IBM Corporation via the CSER Consortium, as well as Icesi University, Cali, Colombia and University of Victoria, British Columbia, Canada.

About the Authors

Norha M. Villegas is a CAS PhD student under the supervision of Prof. Hausi A. Müller. Together with her research groups at University of Victoria (Rigi Research Group - Canada) and Icesi University (DRISO Research Group - Colombia), she is working on the application of software engineering models, techniques and architectures to the development of self-adaptive and self-managing systems. In particular, she concentrates on the implementation of autonomic infrastructures able to support the dynamic adaptation and evolution of context aware applications. Research interests include feedback loops, autonomic computing, dynamic context management, context-awareness, Smart SOA, SOA governance and Smart Internet.

Hausi A. Müller is Associate Dean Research, Faculty of Engineering and Professor of Computer Science, University of Victoria, British Columbia, Canada. He was founding Director of the Bachelor of Software Engineering program. In collaboration with IBM, CA and SEI, his research group investigates methods, models, architectures, and techniques for self-managing and self-adaptive software-intensive systems. Dr. Müller's research interests include software engineering, software evolution, autonomic computing, diagnostics, SOA governance, software architecture, software reverse engineering, re-engineering, program comprehension, visualization, and software engineering tool evaluation. He is Program Co-Chair for IBM CASCON 2010. He was General Chair of VISSOFT 2009 and co-organizer of SEAMS 2006-10 and DEAS 2005. He was Workshops Co-Chair for ICSE 2008. He was General Chair for IWPC 2003 and ICSE 2001. Dr. Müller served twice on the Editorial Board of IEEE Transactions on Software Engineering

(TSE). He is Vice-Chair of IEEE Computer Society Technical Council on Software Engineering (TCSE).

Joanna Ng is currently the Head of Research at IBM Canada Software Laboratories, Center for Advanced Studies. She is also a Senior Technical Staff Member of IBM Software Group. She has held various senior management and architect positions in product development and software strategy within IBM. Joanna is an IBM Master Inventor with a long track record of profitable innovations. She has been granted over twenty five patents from various countries in research areas related to Smart Internet such as mobile commerce, voice-enabled portal, commerce portal, retail industry solutions; service-oriented architecture (SOA); asset repository; semantic and web technologies.

Alex Lau is a research staff member at IBM CAS Canada Research, where his research focuses on service oriented architecture, BPM tools, and business events. Prior to joining CAS, Alex was a software engineer in IBM Canada, where he worked on a number of Websphere BPM products including Websphere Integration Developer and Websphere Business Modeler. Alex holds a Bachelor's degree in Mathematics (Computer Science) and a Master's degree in Mathematics (Computer Science) from the University of Waterloo.

References

[1] T. Berners-Lee, W. Hall, J. A. Hendler, K. O'Hara, N. Shadbolt, and D. J. Weitzner. A framework for web science. *Foundations and Trends in Web Science*, 1(1):1–130, 2006.

[2] J. O. Kephart and D. M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.

[3] R. Krummenacher and T. Strang. Ontology-based context modeling. In *Proceedings Third Workshop on Context-Aware Proactive Systems (CAPS 2007)*, June 2007.

[4] LikeProtocol Working Group. The OpenLike Protocol. http://openlike.org/, 2010.

[5] F. Manola, E. Miller, and B. McBride. RDF Premier. W3C Recommendation. http://www.w3.org/tr/rdf-primer/, 2004.

[6] H. A. Müller, H. M. Kienle, and U. Stege. Autonomic computing: Now you see it, now you don't—design and evolution of autonomic software systems. In A. De Lucia and F. Ferrucci, editors, *International Summer School on Software Engineering (ISSE) 2006-2008*, volume 5413 of *LNCS*, pages 32–54. Springer, 2009.

[7] J. Ng. The Personal Web: Smart Internet for Me. Technical report, IBM Center for Advanced Studies Lab, Toronto, Ontario, Canada, 2010.

[8] J. Ng, M. H. Chignell, and J. R. Cordy. The smart Internet: Transforming the web for the user. Technical report, IBM Canada Center for Advanced Studies, Technical Report, 2009.

[9] A. Pahlevan and H. A. Müller. Static-discovery dynamic-selection (sdds) approach to web service discovery. In *Proceedings of the 2009 World Conference on Services)*, pages 769–772. IEEE Computer Society, 2009.

[10] T. Strang, C. Linnhoff-Popien, and K. Frank. CoOL: a context ontology language to enable contextual interoperability. In *Proceedings 4th IFIP WG6.1 International Conference on Distributed Applications and Interoperable Systems (DAIS 2003)*, volume 2893 of *LNCS*, pages 236–247. Springer, Nov. 2003.

[11] N. M. Villegas and H. A. Müller. Context-driven adaptive monitoring for supporting soa governance. In *Proceedings of the 4th International Workshop on a Research Agenda for Maintenance and Evolution of Service-Oriented Systems (MESOA 2010)*. Carnegie Mellon University Software Engineering Institute, 2010. To appear.

[12] N. M. Villegas and H. A. Müller. Managing dynamic context to optimize smart

interactions and services. In *Proceedings of the Smart Internet Technologies Workshop 2009 (CASCON conference). Springer LNCS*, volume 6400, 2010. To appear.

[13] M. Zuckerberg and B. Taylor. The Open Graph Protocol. http://www.opengraphprotocol.org, 2010.

Context-Aware Personal Web Portals

Jens Weber-Jahnke

University of Victoria

Analytics

&

Visualization

User, Social and Location Context for IBM Business Analytics

Stephan Jou, M.Sc.
Office of the CTO, Business Analytics, IBM Canada

Gartner, in their Gartner Hype Cycle for Context-Aware Computing 2010, provides the following statement: “Context-aware computing is about improving the user experience for customers, business partners and employees by using the information about a person or object's environment, activities, connections and preferences to anticipate the user's needs and proactively serve up the most appropriate content, product or service. Enterprises can leverage context-aware computing to target prospects better, increase customer intimacy and enhance associate productivity and collaboration.”

Data availability is no longer an issue for Business Analytics applications. Now the challenge is to deliver the right information at the right time in a way that is relevant and understandable.

Without context more data from more sensors means ever growing piles of puzzle pieces. Harnessing the information explosion necessitates placing information into context. In a sense, the data must find the data before relevance/insight/action can be taken with confidence.

IBM Business Analytics started with support for location context when IBM Cognos 8.4 Go! Mobile shipped with location awareness built in (you can auto-filter your reports by your current location, as read in from your BlackBerry GPS). We also have a rich metadata repository in our models which we have begun to tap in recent versions of Go! Search and other parts of our portfolio. For example, in Go! Search 8.4 you can enter in a query like "Top sales for Tents in Japan", and -- if such a report does not exist which contains those terms, it will use the metadata that we have around the customer's models and dynamically construct a query or report that answers the question.

By focusing on support for embedded context information within the Business Analytics platform and other products areas, IBM is well positioned to provide our customers with scalable context-aware solutions in areas such as predictive analytics, search, content analytics, real-time monitoring, decision management, and many more.

As part of Smarter Planet and its related strategies, many parts of IBM are investing in technologies that enable us to provide context aware computing. This includes ongoing investments in linguistic analysis, social networking, semantics and relevance, etc. More specifically, there are specific strategic investments that can be made in the areas of user, social and location context.

User context, closely tied to the personal web, recognizes the fact that each user of our software is different and that our complex software can be made much more accessible by personalizing. Knowledge of who a user is, built up based on observations on user behaviors and stated preferences, can be used to provide appropriate subsets and customizations of the interface, or even anticipate the user's requirements and needs. If a user runs the same report every morning at 10am, could our software not anticipate that and provide the report in advance?

Social context recognizes that a user does not operate as a single entity, but rather as part of a larger group of people in both formal and informal networks, and that these networks can be harvested for information that can provide value back to the user. If a user's boss finds great value in a specific financial report, then perhaps the user should too; if many of the user's peers express concern in a predicted budgetary figure, then perhaps the user should be alerted to a lack of trust in that figure on his or her report.

Location context goes beyond who the user is, extending into the realm of where the user is. The previous example of location-based filtering can be extended into more sophisticated forms of location intelligence within Business Analytics applications. Can a report be filtered not by the simple geometrical area around the user's current location, but rather by the region prescribed by less than 15 minutes drive distance from the user's current location, at the current time of day?

Information without context is akin to evaluating a stand-alone puzzle piece. Organizations that leverage information 'in context' will be substantially more competitive and they will also recognize opportunity and mitigate risk with new levels of prediction accuracy.

Context aware computing will revolutionize the entire spectrum on business analytics ranging from mobile, search, customer insight to fraud detection and advances in predictive analytics. We will continue to see context rich analytic capabilities integrated across IBM's portfolio, and our customers will use these advances in prediction to be more efficient and deliver high quality services to their customers.

Towards Smart Visual Data Exploration Tools

Lars Grammel and Margaret-Anne Storey

University of Victoria

Abstract

The semantic information made available by the Smart Internet could enable visual data exploration tools to provide users with a personalized and task-centric data analysis experience. We describe some of the potential opportunities and challenges of leveraging user profiles, task context, and semantic meta-data to improve visual data exploration and analysis.

1 Motivation

Visual data exploration and analysis is a time-consuming activity that requires both domain specific knowledge (about the data being explored) and analytical expertise (in both visualizations and data analysis). It is often performed in the context of a higher-level task. When visual data exploration is carried out by a single analyst, this process would ideally be tailored to the individual, to the data being explored, and to the contextual task. The Smart Internet [3], "a platform for automatic, dynamic aggregation of data and services for the purpose of supporting each user's goals, tasks and concerns, both cognitively and socially" [2], could provide such semantic information that would enable streamlining this activity and that would make visual data exploration tools accessible to a larger audience.

In this position statement, we describe ideas of how user profiles, task context, and semantic meta-data could potentially be employed to create adaptive visual data exploration tools, and some of the challenges that still need to be overcome[1]. Smart visual data exploration tools that adopt to the user and his/her task could be used as components of the Personal Web [2] that provide data analysis capabilities as required by the user.

We will first briefly summarize the gap that exist between generic visual data exploration tools and custom designed visualizations, and examine at related research that addresses this gap. Next, we will describe potential opportunities and challenges of leveraging user profiles, task context and semantic meta-data. Finally, we describe Choosel, a visual data exploration tool we have developed, and the adoption capabilities we are planning to add to it.

2 Related Work

Visual data exploration tools allow users to analyze and gain insights into data sets visually. Commercial tools such as IBM Cognos 8 Business Intelligence[2], Tableau[3], and Tibco Spotfire[4] provide standard visualization and analysis tools that work with diverse sets of data. These differ from custom-designed, data-specific visualizations in that they can be applied out of the box (eliminating the delay required for developing a domain-specific visualization). However, the resulting visualizations are inherently oblivious of the data set, the user and the task context and thus

[1] The discussion is limited to desktop computers, ignoring device-specific adaption.
[2] http://www-01.ibm.com/software/data/cognos/products/cognos-8-business-intelligence/
[3] http://www.tableausoftware.com/
[4] http://spotfire.tibco.com/

are potentially less effective than custom made visualizations (which are designed for a specific data set and for a typical user group and task).

This gap between visual data exploration tools and custom designed visualizations is partially addressed by automatic visualization tools. Automatic visualization is concerned with creating visualizations without the involvement of a visualization designer. Automatic visualization tools such as APT [9] usually take a data set and a list of data properties as input. Some tools attempt to leverage task descriptions [1, 11], however, tasks are usually hard to describe and the visualization intent may be difficult to elicit [10]. Automatic visualization functionality has also been integrated into the user interface (UI) of visual data exploration tools such as Tableau [10]. Gilson *et al.* employed semantic domain descriptions for automatic visualization using domain, bridging, and visualization ontologies [4]. The domain ontology described the problem domain, the visualization ontology described a specific visualization, and the bridging ontology contained expert knowledge on how the two should be mapped. Despite the progress researchers made in automatic visualization, leveraging task context, user profile and semantic meta-data is still require further investigation.

3 Towards Smart Visual Data Exploration Tools

In this section, we describe how a hypothetical visual data exploration tool could adopt itself based on user profiles, task context and semantic meta-data, and what challenges we foresee. In our scenario, the data exploration tool would rely on the Smart Internet to provide the user profile and task context. A data warehouse would provide the data as well as semantic meta-data.

3.1 User Profile

Four adaption methods that could be derived from the user profile are preference-based adaption, impairment-based adaption, skill-based adaption, and social-network-based adaption.

Preference-based adaption is based on user preferences, such as language settings or regional preferences (e.g., time zone or currency). Impairment-based adaption takes the impairments of a specific user into account (e.g. color deficiencies). Skill-based adaption takes the user's skill level into account to provide support and hide/offer functionality. Social-network based adaption considers potential support in the social network during UI customization.

The following is an example of using skill-based and social-network based adaption. User Carl is analyzing some business data of interest. He selects several data attributes he wants to analyze. Because he is inexperienced in creating visualizations, the system decides to present the most appropriate visualization, instead of letting him choose from a range of potential visualizations. While Carl tries to analyze the visualization, the system detects that Carl performs repetitive operations and seems to be stuck. Depending on the level of expertise and availability of people in Carl's company's social network, the system could either recommend a local expert for Carl to contact, or display a help page on how to interpret such a visualization.

Challenges to employing user profile-based adaption methods include how to determine the skill level of a user, detecting skill level changes, knowing when a user is stuck, and how to support the user in a non-intrusive manner.

3.2 Task Context

The task context models the high-level task that the user is working on. However, there are several levels how tasks and activities can be modeled and are connected [5]. While the Smart Internet might be able to provide information about the higher level task, knowing concrete intentions for using visualizations at a lower task level might still be challenging [10].

The task context can be used to select relevant data from the data warehouse, to guide automatic visualization algorithms, and to find visualizations used previously by others on similar tasks. Challenges are the modeling of the task, and how it can be connected to automatic visualization algorithms.

3.3 Semantic Meta-Data

Semantic meta-data describes how data and data attributes in the data warehouse relate to real-world processes and concepts. This includes relationships to concepts such as units, relationships to other data attributes, and formulas how data values are calculated. For example, semantic meta-data could specify that the budget sales data attribute refers to a projection for future sales in US dollars. The general concept of sales could then link to other data attributes with the same underlying type, such as the real sales, or to other sales estimations.

Similar to the task context, semantic meta-data can be used to direct access to parts of the data warehouse and to inform automatic visualization algorithms [7]. It can also be used to suggest other possible data attributes to visualize or to explain observed relationships in the data.

For example, when the user requests to visualize actual and estimated sales, the automatic visualization algorithm could take into account the relationship between those two data attributes. It could then choose a visualization that highlights the differences between actual and estimated sales for past months.

We are planning to explore different ways of leveraging such semantic meta-data as part of our development of the visual data exploration tool Choosel.

4 Choosel

We have started developing a visual data exploration tool called Choosel[5] that will enable us to investigate some of these concepts [6]. Choosel is a web-based tool that aims at facilitating flexible visual data exploration for information visualization novices (see Figure 1). It supports the iterative construction of multiple coordinated views during the visual data analysis process.

We have built a work item exploration tool (Work Item Explorer) and a bio-medical data exploration tool (Bio-Mixer[6]), which are based on the Choosel framework. The Work Item Explorer is an example of our vision of software mashups [8] and allows developers to analyze work item data in different visualization such as graphs, time lines, and charts. An online version of Choosel is also available to enable web users to explore any kind of supplied tabular data[7].

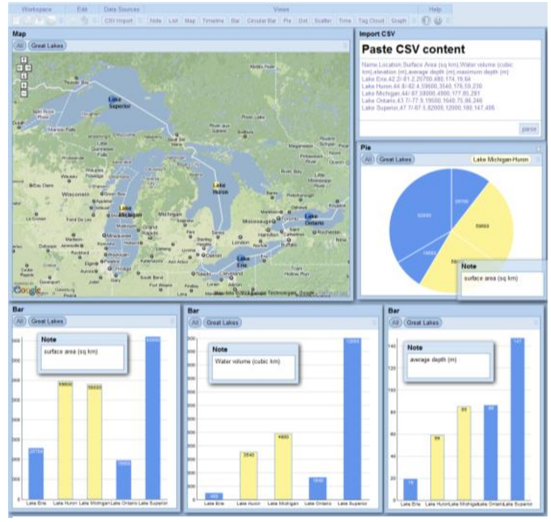

Figure 1: Choosel Visual Data Exploration Environment

We are currently working on adding automatic visualization and visualization recommendation to Choosel. As part of this, we are specifically looking at using data sets that are annotated with semantic meta-data.

5 Conclusion

Smart visual data exploration tools have the potential to make visual data exploration more effective and accessible for inexperienced users. Leveraging user profiles, task context and semantic meta-data are fruitful avenues for adopting such tools and directing automatic visualization tools. We believe that smart visual data exploration tools will be an important part of the Smart Internet and enable users to rapidly understand relevant data.

[5] http://code.google.com/p/choosel/
[6] http://bio-mixer.appspot.com/
[7] http://choosel-mashups.appspot.com/

Acknowledgements

The authors would like to thank Chris Bennett for his editing support. This research was funded by an IBM CAS PhD Fellowship.

References

[1] Stephen M. Casner. Task-analytic approach to the automated design of graphic presentations. *ACM Transactions on Graphics (TOG)*, 10(2):111–151, 1991.

[2] Mark Chignell, James R. Cordy, Joanna W. Ng, and Yelena Yesha. First symposium on the personal web. In *CASCON '10: Proceedings of the 2010 Conference of the Center for Advanced Studies on Collaborative Research*, 2010.

[3] Mark Chignell, James R. Cordy, Joanna W. Ng, and Yelena Yesha, editors. *The Smart Internet*, volume 6400 of *Lecture Notes in Computer Science*. Springer, 1st edition, 2010.

[4] Owen Gilson, Nuno Silva, Phil W. Grant, and Min Chen. From web data to visualization via ontology mapping. In *Computer Graphics Forum*, volume 27, pages 959–966. Blackwell Publishing Ltd, 2008.

[5] David Gotz and Michelle X. Zhou. Characterizing users' visual analytic activity for insight provenance. *Information Visualization*, 8(1):42–55, 2009.

[6] Lars Grammel and Margaret-Anne Storey. Poster: Choosel - web-based visualization construction and coordination for information visualization novices. In *InfoVis 2010: Proceedings of the IEEE Symposium on Information Visualization*. IEEE Computer Society.

[7] Lars Grammel, Melanie Tory, and Margaret-Anne Storey. How information visualization novices construct visualizations. *IEEE Transactions on Visualization and Computer Graphics (TVCG, Proc. InfoVis 2010)*, 16(6), 2010.

[8] Lars Grammel, Christoph Treude, and Margaret-Anne Storey. Mashup environments in software engineering. In *Web2SE '10: Proceedings of the 1st Workshop on Web 2.0 for Software Engineering*, pages 24–25, New York, NY, USA, 2010. ACM.

[9] Jock Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics (TOG)*, 5(2):110–141, 1986.

[10] Jock Mackinlay, Pat Hanrahan, and Chris Stolte. Show Me: Automatic presentation for visual analysis. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1137–1144, 2007.

[11] Steven F. Roth and Joe Mattis. Automating the presentation of information. In *Seventh IEEE Conference on Artificial Intelligence Applications, 1991.*, 1991.

Supplementing Semantics with Statistics in the Personal Web

Mark H. Chignell and M. Ryan Kealey

Interactive Media Lab, Department of Mechanical and Industrial Engineering,

University of Toronto

1 Introduction

In her position paper for the First Symposium on the Personal Web", Ng (2010) describes the Personal Web as "user-sovereign web integration" resulting in the generation of dynamic and highly personalized web artifacts for visualizations and interactions. She envisions a process of "normalization and abstraction" that will converts entities on Web pages into integrate-able web elements suitable for Personal Web interactions. A "simple, well-defined meta-model" would then be used to generate visualizations and interactions based on the web integration results.

We will refer to Ng's vision of the personal web as "Plan A". Plan A assumes sufficiently well described Web content in the form of RDF annotation, and a sufficiently powerful semantics that can reason about, and link together, content from a variety of Websites. In this paper we propose a statistical approach as a kind of "Plan B" that may, in some situations, still provide useful personal web functionality in cases where Plan A fails.

2 Background

In his roadmap of the semantic web Berners-Lee (1998) characterized RDF as a necessary first step in regularizing semantics across websites. Ng (2010) posited that users could "gather personal linked data RDF graphs from several web domains and group them together into one semantic context". Berners-Lee saw an underlying (predicate) logic layer that would supplement RDF assertions with a powerful linking and querying capability, suggesting that different inference engines might be developed for different applications.

One view of the Personal Web is that it a personalized and contextualized version of the semantic web, motivated by the need to create smart interactions that are personalized according to the contexts and interests of individual users. Seen in this way, the Personal Web should be a more tractable version of the Semantic Web. The semantics of online shopping, travel, entertainment, self improvement and education, and personal health etc., are highly constrained relative to semantics in general. In this contextualized view, specialized inference engines make sense, perhaps forming a collection of expert systems (cf., Parsaye and Chignell, 1988) that provide the reasoning capabilities required by the Personal Web. However, while expert systems have had some success in specialized areas, they tend to be brittle, in that they fail to reason correctly when assumptions that they were built on don't hold and they can't fall back on the common sense knowledge that people tend to use when they don't have required expertise to deal with a problem.

The CYC project (Lenat and Guha, 1990) aimed at capturing "common sense knowledge" so that reasoning could be more generalized and less brittle. However, experience in artificial intelligence over recent decades suggests that it is much easier to develop reasoning engines than it is to represent the knowledge that is reasoned with. Generalized reasoning of the type required for natural language understanding has yet to be achieved.

In summary, while the prospects for developing a semantic personal web are likely better than the semantic web as a whole, it is still a highly challenging task, since people and contexts vary along many different dimensions creating many different types of knowledge that need to be represented and reasoned with.

In natural language understanding, statistical reasoning has been used to supplement knowledge-based reasoning. With respect to processing of speech, Callison-Burch and Osborne (2003) claimed that: "Statistical techniques in speech recognition have so vastly outstripped the performance of their non-statistical counterparts that rule-based speech recognition systems are essentially no longer an area of research."

Sheth et al (2005) noted that exploiting heterogeneous data in the semantic web would require a broad range of semantics, which they classified into three forms: implicit; formal; powerful. In their characterization of "powerful" semantics: "Statistical techniques give us great insight into a corpus of documents or a large collection of data in general…. All derived relationships are statistical in nature and we only have an idea or a likelihood of their validity."

3 A Statistical Approach

In this paper we propose that explicit methods for representing semantics in the Personal Web through representation methods such as RDF and reasoning methods such as predicate logic be supplemented with statistical analysis (data analytics) and interactive visualization. The basic idea is that, in cases where explicit semantic reasoning cannot identify a sufficient set of integrated web elements that address the users need, interactive visualizations based on focused statistical analyses may provide users with "sufficiently convenient" overviews of relevant content and actions.

Appropriate information visualization has been proposed as a way to replace effortful thinking with simpler and more direct seeing (Card et al., 1999). A powerful demonstration of this idea, with direct manipulation of sliders to interactively manipulate the visualization (through inferred querying to the underlying database) was provided in the dynamic querying project (Williamson and Shneiderman, 1992). Since then there have been many techniques developed for interactive visualization (e.g., the elastic hierarchies developed by Zhao et al., 2005). New toolkits have been developed that greatly simplify the task of building interactive visualizations (e.g., Prefuse, Heer et al., 2005). These toolkits enable the development of applications that construct interactive visualizations to perform novel tasks such as tracking the popularity of baby names over time using census data (www.babynamewizard.com). Visualization toolkits are also leading to new forms of collaborative visualization, where data may be explored, interesting visualizations discovered, and then annotated and shared as part of larger discussions within blogs and communities of users (Heer et al, 2009).

The work of Casner (1991) and Wilkinson (2005) provide detailed perspectives on, and develop the idea of, task-based graphical presentation. In the proposed system, automatic analysis of the dataset would follow from analysis of the user's context and current tasks. Applications or 'task-based overlays' could then be developed on top of this system for facilitating the work of different different people and occupations, and the tasks associated with them. For instance, nurses in critical care units might use analytics-driven interactive visualization to keep better track of in-need patients (potentially reducing the incidence of "failure to rescue"). In a related example, emergency physicians could monitor the status of their other patients while they work with a particular patient. Continuing the healthcare example, but with a different occupational role, hospital administrators could monitor the overall situation and use general patterns of patient status to provide better measures of future bed requirements and availability, as well as identifying resource bottlenecks and the like.

4 A Healthcare Example

Healthcare is a domain where there is a huge amount of data, and where the available semantics and knowledge tend to apply to populations rather than individuals. Supplementing general medical knowledge with relevant views of how relevant peers of the current case respond to various treatments may provide improved clinical decision support. The focus in this approach would be on the analytic and visualization tools that would take large volumes of (ideally stored in the form detailed electronic health records that include real-time data) healthcare data in the repository and make it available in the form of summaries and interactive visualizations. These tools would include parameters that could be set to allow healthcare professionals to customize their summarized view of the data. Algorithms could be built that would automatically generate visualizations of data depending on the context of the user and the dataset currently under review. Examples of automated statistical analyses (based on data patterns) that could generate appropriate visualizations were presented by Dan Rope at IBM University Day (Markham Ontario, April 2010). Making this approach work in realistic applications will likely involve extensive analyses of user requirements, followed by the prototyping and refinement of application user interfaces that address those requirements. The resulting application user interfaces then serve as specifications of what analytics and visualization capabilities are required in order to populate the user interface appropriately. We envisage the development process as a three legged stool where the work is supported by:

- Requirements analysis and user cases concerning the applications needed by healthcare professionals.
- Development of an analytics and visualization framework for automatically showing events, patterns and trends in large amounts of healthcare data.
- Prototyping and user testing of healthcare applications built using the analytics and visualization framework

The analytics and visualization framework will influence what can be prototyped, but at the same time knowledge of application requirements and user interface features may influence the functionality that is developed for the analytics and visualization framework. Similarly, application requirements will create a space of required functionality for visualization and analytics, but at the same time the analytics and visualizations that are developed may inspire and constrain application requirements.

For the analytics portion of the work we propose to utilize the SPSS statistical package, writing procedures that automate the use of particular techniques (such as time series analysis, cluster analysis and regression) to permit the conversion of the data into interpretable patterns, events, and trends. If possible we may build on the work that IBM/SPSS have already been doing in this area. In addition we propose to use Cognos tools to to organize data and to provide visual summaries of the key patterns and trends.

5 Conclusions

The Personal Web is a vision that requires identification and linking of integrateable Web entities. The success of Personal Web applications will depend to a large extent on the adequacy of these personalized and contextualized identification and linking operations. Characterizing Web semantics, and semantics in general, has proven to be a challenging problem and it remains to be seen how rapidly effective approaches to Personal Web Semantics will evolve. Personal Web statistics is presented as a complementary approach that will provide more flexibility in dealing with large, idiosyncratic, and poorly indexed datasets at the cost of requiring more interaction and exploration of the user. In many domains where users are highly motivated to achieve their tasks goals, this would seem to be a reasonable tradeoff.

Acknowledgements

The ideas expressed here have been influenced by a number of researchers ranging from Jock Mackinlay and others to Lars Grammel, Stephan Jou and Dan Rope. Much of our thinking on clinical decision support has grown out of earlier research with Sharon Straus. The research was supported by an IBM Faculty Award to the first author.

References

[1] Berners-Lee, T., Hendler, J. & Lassila, O. (2001). The Semantic Web. Scientific American, 284(5), 34-43.

[2] Callison-Burch, C. and Osborne, M. (2003). Statistical Natural Language Processing. In A. Farghaly (Ed.), Handbook for Language Engineers. CSLI Lecture Notes, University of Chicago Press.

[3] Card, S. K. , Mackinlay, J. D. and Shneiderman, B. (Eds., 1999). *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc.

[4] Casner, S. M. 1991. Task-analytic approach to the automated design of graphic presentations. *ACM Trans. Graph*. 10, 2 (Apr. 1991), 111-151.

[5] Heer, J. (2005). Prefuse: a Toolkit for Interactive Information Visualization. In CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems, pages 421—430. NY: ACM Press.

[6] Heer, J., Viégas, F. B., and Wattenberg, M. 2009. Voyagers and voyeurs: Supporting asynchronous collaborative visualization. *Commun. ACM* 52, 1 (Jan. 2009), 87-97.

[7] Lenat, D. and Guha, R. V. (1990). Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project. Addison-Wesley.

[8] Parsaye, K. and Chignell, M.H. "Expert Systems for Experts", New York, John Wiley & Sons, 1988.

[9] Sheth, A., Ramakrishnan, C., & Thomas, C. (2005), "Semantics for the Semantic Web: The implicit, the formal and the powerful," *International Journal on Semantic Web & Information Systems*, vol. 1, no. 1, pp. 1–18.

[10] Streitz, N. (1987). Cognitive compatibility as a central issue in human-computer interaction: Theoretical framework and empirical findings. In: G. Salvendy (Ed.), *Cognitive engineering in the design of human-computer interaction and expert systems*. Amsterdam: Elsevier Science. 75 - 82.

[11] Takeshita H, Davis D, Straus S. Clinical evidence at the point of care in acute medicine: a handheld usability case study. *Proceedings of the human factors and ergonomics society 46th annual meeting*, 2002, p. 1409-13.

[12] Wilkinson, L. (2005) *The Grammar of Graphics, 2nd Ed*. New York: Springer- Verlag.

[13] Williamson, C., and Shneiderman, B. (1992). The Dynamic HomeFinder: Evaluating dynamic queries in a real-estate information exploration system, Proc. ACM SIGIR'92 Conference, Copenhagen, Denmark, (June 1992), 338-346.

[14] Zhao, S., McGuffin, M.J., and Chignell, M.H. "Elastic Hierarchies: Combining Treemaps and Node-Link Diagrams," Information Visualization, IEEE Symposium on, p. 8, 2005 IEEE Symposium on Information Visualization (InfoVis 2005), 2005

Healthcare

Information Feudalism: How legacy information systems kill people in the healthcare system

Jacques Lee
Sunnybrook Health Sciences Centre

Unfortunately, every day in the current health care system in Canada, health care professionals must treat patients while lacking critical information needed to provide them with optimal care.[1] Here is an example from a recent patient encounter. The names of the individuals, and the colours of the pills involved have been changed to protect the innocent.

“Mrs. Smith, the reason you’ve been feeling so rotten is because you have a pneumonia in your right lung. We’d really like to start treatment right away, but you told the nurse you had a serious reaction to an antibiotic before. Do you remember which one?”

Mrs. Smith replied “The Pink one.” Sensing that this was not particularly helpful she shot back “It’s all written down on my medical record – can’t you just look it up?”

Indeed, Mrs. Smith’s assertion seem reasonable enough. If the internet guarantee’s that I find out about a toddler smoking in a small village in the Philippines, surely it is not unreasonable to think that I should be able to find out about a significant adverse event that occurred at another institution within a single-payer, publicly funded healthcare system.

What Mrs. Smith doesn’t know is that not only is it difficult to get that information from across town, I often can’t get critical information like that from down the hall in my own hospital. Something is rotten in the state of healthcare information systems.

[1] Stiell A, Forster AJ, Stiell IG, van Walraven C. Prevalence of information gaps in the emergency department and the effect on patient outcomes. *Cmaj.* Nov 11 2003;169(10):1023-1028.

The Personal Web and Clinical Medicine: A Physician's View

Tammy Sieminowski MD MEng CCFP
Bridgepoint Hospital
Toronto, Ontario

Abstract

This paper will discuss the Personal Web as it pertains to supporting the practice of clinical medicine. An overview of the practice of clinical medicine is provided. Current challenges stem from heterogeneous data collection and storage formats; insufficient links between existing electronic resources, and lack of smart services and interactions to support the use of this information both at the point of care and across clinical practices. Potential research directions envision the Personal Web supporting clinical medicine with the ultimate goal of improved patient care through patient-centered data aggregation and integration as well as smart interactions and smart services that support the use of evidence based decision tools and assist with practice management.

1. Introduction

The Personal Web has been described as the user view of the smart internet. The smart internet has been defined as being composed of smart interactions and smart services. Smart interactions pertain to the discovery, aggregation, and delivery of data. Smart services are comprised of data infrastructure and task management [1]. Currently, electronic based health care data and resources are increasing but not ubiquitous. Links are forming but networks are incomplete.

Smart interactions and smart services are scarce. Data aggregation, transmission, and analysis are still largely manual. An environment that supports context specific user driven integrations across the web is desirable. The Personal Web holds great promise in supporting the practice of clinical medicine, with the ultimate goal of improved patient care. This paper will provide an overview of the practice of clinical medicine, discuss challenges, and potential research directions.

2.0 The Practice of Clinical Medicine

The practice of clinical medicine can be divided into the following categories: clinical decision making, requests for further services, documentation, practice management, and ongoing education. This classification does not encompass procedures.

2.1 Clinical Decision Making

The process of clinical decision making requires the collection of relevant patient specific data in a timely fashion. The time scale is in the order of minutes. The absence of information often cannot delay decision making. Information is gathered from the patient, involved caregivers, other clinical team members, other involved physicians, documented past medical history, laboratory results, imaging results, and decision support tools such as guidelines, reference material, risk calculators, etc. Formats include: direct or indirect verbal communication, paper based records and communications, electronic records. Sorting relevant from irrelevant information ('signal from the noise') is one of

.

the main challenges in clinical decision making. More information is not always better.

As information is gathered, the physician concurrently formulates the problem and generates possible solutions. These critical calculations are constrained by cognitive capacity. Currently, no analytic tools exist to generate final optimal recommendations given a patient's unique parameters.

2.2 Requests for further services

Having gathered available relevant information, management of a patient's issues may require requests for further services**.** This may include requests to obtain further patient specific information such as 'labs' (analysis of bodily fluids) or imaging (such as CT scans, ultrasounds, etc). Sometimes, the opinion of another physician is required ('referral') or requests are made for specific procedures. Each of the foregoing requires the requesting physician to specify the service and forward selected and/or required patient information to assist the next provider in the interpretation of clinical picture, urgency, and results. Transmission of this information varies according to the recipient's permissible formats.

2.3 Documentation

Physicians are not only readers of the patient record, they are also the authors. Documentation is required for a number of reasons: to contribute to the local patient record for the purposes of personal recollection and thought organization; to provide communication to other health care providers who will view the patient record in order to gather information necessary for their decision making around patient issues; to serve as an historical record of the events that transpired (particularly in the event of undesirable outcomes); to support financial remuneration.

2.4 Practice management

Most physicians practicing clinical medicine manage a number of patients at one time. This may vary from a few patients in an ICU setting, to several dozen during a shift in an emergency department, to a few hundred for specialists, or upwards of 2500 patients for a family physician. Beyond seeing one patient at a time, physicians must manage incoming information for all of their patients. Information arrives irregularly and by various means: mail, fax, phone messages, direct verbal communication from other health team members, electronic health record, email, etc. New information must be managed in terms of urgency. All information must be reviewed and acknowledged in some manor. Management of this information is usually up to the individual physician, subject to the influence of their environment (e.g. group vs. solo vs. hospital practice).

To date, support of overall practice management has been limited. The Ministry of Health in Ontario has started to provide subgroups of physicians with some practice specific clinical audit information, such as the percentage of their patients who have received an indicated screening test. Some electronic health records may support scheduling, information management and reconciliation, clinical audits, etc. Incorporating patient information from multiple sources (in heterogeneous formats) into a specific EHR to allow for this analysis remains a challenge.

2.5 Ongoing Education

Physicians require up to date information in two general circumstances: in a given clinical context or as a general knowledge base update. The latter is often referred to as Continuing Medical Education or 'CME'. Information for a given clinical context is usually required instantaneously and may be obtained via database search, website specific searches, general search engines (e.g. Google), paper based references, or colleague query. Electronic information is usually accessed through mobile or desktop devices. Finding the answer to a specific question is often not straightforward.

The process of formal CME is varied and an ongoing area of research [2]. The overall goal of CME is to improve the quality of healthcare. Whether or not this is actually realized is unclear. CME includes but is not limited to: journal reading (paper or electronic), medical meetings or educational rounds, academic detailing, informal discussions amongst colleagues, self study, individual practice audits, public media stories, etc. These events may be scheduled, prompted (arrival by mail, email, RSS feed, etc), or occur spontaneously.

The amount of new and existing research based clinical data is immense. The National Institutes of Health's Medline® indexes 5485 journals [3]. In an effort to assist in the management of this massive amount of information, clinical guidelines and systematic reviews have been developed however, often more than one set of guidelines will exist on a particular topic. Repositories for these resources are in turn immense: The National Guideline Clearinghouse currently houses more than 7000 guidelines [4] and The Cochrane Library more than 4000 systematic reviews [5].

Access to information varies. Some journals, such as the British Medical Journal, update their publications daily and provide open access to all articles [6]. Others, such as the New England Journal of Medicine, publish weekly and limits access to full articles only to subscribers [7]. Direct links between electronic health records and these evidenced based resources are lacking.

3.0 Current challenges

The practice of clinical medicine today faces a number of challenges. These may be summarized as heterogeneous data collection and storage formats, lack of links between existing electronic resources, and lack of smart services and interactions to support the use of this information. Functionally, the problems may be divided as follows:

3.1 Access to existing information

One of the greatest challenges facing physicians in Ontario today is the timely access to existing documented patient information. Access is limited by heterogeneous, agency-centric data storage methods and privacy concerns. For the purposes of this paper, 'agency' will include: hospitals, laboratories, community pharmacies, and physician offices. Within an agency, each patient may have a paper as well as an electronic chart, and one software system may not be linked to another requiring the user to 'exit' one system and 'enter' another to access information on the same patient. External linkages between health care agencies are forming but remain scarce. Searching for relevant data within an existing electronic record is often a challenge as user interfaces are inflexible, search features are lacking or absent, and software architecture varies from vendor to vendor.

Individual agencies are strictly charged with the privacy and protection of patient information. If information housed in another agency is required by a physician, current policies and procedures around these important privacy concerns often necessitate permission to be obtained on a case by case basis before the information can be released, something that is often not feasible in a timely manor especially in urgent situations or outside of routine business hours. Some links between participating agencies are starting to form, facilitating view only access to existing patient data.

Accessing up to date, relevant data to support evidence based point of care decision making is another challenge. As discussed, the body of literature is immense. Though many clinical guidelines exist, studies have found adherence rates to be as low as 27% [8]. Factors influencing uptake include ease of use and time constrained work environments. Cumbersome access and inefficient search methodologies may be contributing to these issues.

3.2 Dissemination of information

Not dissimilar to the problems associated with accessing information, information dissemination is plagued by heterogeneous formats and requirements. Documentation may occur in a paper and/or electronic chart; orders may be written or electronically entered; requisition forms are paper based or electronic; most referrals are handwritten and faxed to the next office for review; prescriptions may be hand written or electronically generated, and are then transmitted by hand or fax. Requirements are agency specific and electronic links are forming but still lacking.

3.3 Poor uptake of electronic health records

In a 2008 survey of Canadian hospitals, 54% were identified as having some sort of EHR. Fewer than 3% had no paper record [9]. Electronic health records have not performed as anticipated. There has not been a clear reduction in errors and cost, nor an increase in provider productivity [10]. Documentation time has increased [11]. Undesired effects, such as increased mortality, have been observed [12]. Reasons for poor performance may include: cost, implementation and integration issues, and poor usability. Software is inflexible and often not well suited individual user's needs, forcing users to alter workflow to accommodate software requirements. Service aggregation [13], where by users decide select and integrate services is lacking.

Without electronic records, we have no potential to connect and integrate them.

3.4 Information and Practice Management

As discussed in 2.4, physicians must manage information across the spectrum of their patients. On a daily basis, this means accumulating incoming information, reviewing each piece, and triaging accordingly. Information may be categorized as 'no further action required', or 'action required'. If no further action is required, that information may be 'filed' in the patient chart. If further action is required, decisions must be taken regarding level of urgency, etc. Incoming information arrives sporactically around the clock via phone, fax, mail, email, within an EHR, or via handwritten notes. Currently, it is up to the individual physician to develop a plan to accumulate and manage all of this information and the resulting actions, and to develop some sort of reconciliation strategy. Electronic services supporting integration and management of information across an entire practice are lacking. Similarly, there is little or no automated data collection, analysis, and feedback related to general practice management issues such as scheduling optimization, etc.

4.0 Research directions

The Personal Web holds great potential to support the practice of clinical medicine with the ultimate goal of improved patient care. This may be measured in terms of clinical outcomes, error reduction, patient satisfaction, etc. The 'user' may be viewed from the patient or physician perspective. Following are proposed research questions:

4.1 Patient-centric data

Can smart interactions and services support the collection and aggregation of patient information from multiple agencies, for multiple users? The Personal Web would allow the creation of patient based sites that would automatically aggregate laboratory results, imaging results, medication records, documentation from involved physicians, hospitalizations, etc. Patients would access and contribute to their own records while the integrity of other information is maintained. Interfaces would be user specific and flexible; defined by need and context; ensure privacy; support quick search and aggregation of relevant web based information and services.

4.2 Practice management

Can smart interactions and services support the aggregation and management of data as it is generated across a clinical practice? The Personal Web would allow the physician-user to easily automate practice-wide aggregation of incoming information; customize alerts; bin and batch according to review preference; automate reconciliation; provide analysis and feedback to optimize efficiency; support smart interactions pertaining to tasks that require further action.

4.3 Decision Support and Continuing Medical Education

Can smart interactions and services facilitate the delivery and discovery of context specific information to support point of care clinical decision making and ongoing education? The Personal Web would allow the physician user to easily and directly access context specific information on demand; provide automated prompts; allow for customized information delivery.

5.0 Summary

The practice of clinical medicine is complex. Current challenges in accessing and disseminating information stem from heterogeneous data collection and storage methods; insufficient links between existing electronic patient records; lack of automated data aggregation, integration, and utilization; inefficient access to an immense body of existing evidence based research. Built on smart interactions and smart services, the Personal Web holds great potential to facilitate the practice of clinical medicine, with the ultimate goal of improving patient care.

About the author

Tammy Sieminowski is a lecturer in the Faculty of Medicine at the University of Toronto and attending physician in Neurorehabilitation at Bridgepoint Hospital. She has worked in community based family practice offices, emergency departments, and inpatient settings in Canada, Australia, and the United States. She received her medical degree from the University of Toronto, where she also did her post graduate training in Family and Community Medicine. She recently completed a Master of Engineering degree in the Department of Mechanical and Industrial Engineering at the University of Toronto. Her research interests include forecasting length of stay, healthcare based process improvement and optimization, and physician workflow.

References

[1] Chignell M, Cordy JR, Ng JW, Yesha Y. First Symposium on the Personal Web**.** CASCON Proceedings, November 2010

[2] Nylenna M, Aasland OG. Doctors' learning habits: CME Activities among Norwegian physicians over the last decade. *BMC Medical Education*, May 2007

[3] ://www.nlm.nih.gov/bsd/num_titles.html Accessed 3 Oct 2010

[4] ://www.guideline.gov/index.aspx Accessed 3 Oct 2010

[5] ://www.cochrane.org/about-us Accessed 3 Oct 2010

[6] ://resources.bmj.com/bmj/about-bmj Accessed 3 Oct 2010

[7] ://www.nejm.org/ Accessed 3 Oct 2010

[8] Francke AL, Smit, MC, de Veer AJE, Mistiaen P. Factors influencing the implementation of clinical guidelines for health care professionals: A systematic meta-review. *BMC Medical Informatics and Decision Making*, September 2008

[9] Urowitz S, Wiljer D, Apatu E. Is Canada ready for patient accessible electronic health records? A national scan. *BMC Medical Infomatics and Decision Making,* July 2008

[10] Sidorov J. It Ain't Necessarily So: The Electronic Health Record And The Unlikely Prospect Of Reducing Health Care Costs**.** *Health Affairs*, pages1079 – 1085, 2006

[11] Poissant L, Pereira J, Tamblyn R, Kawasumi Y. The Impact of Electronic Health Records on Time Efficiency of Physicians and Nurses: A Systemic Review**.** *Journal of the American Medical Informatics Association***,** October 2005

[12] Han YY, Carcillo JA, Venkataraman ST. Unexpected Increased Mortality After Implementation of a Commercially Sold Computerized Physician Order Entry System. *Pediatrics*, pages 1506-1512, December 2005

[13] Czarnecki K. Modeling Service Aggregation for the Smart Internet: Challenges Ahead. Position paper for CAS Workshop on Smart Internett Technologies at CASCAN '09.

Smart Interactions for Health Self Management

Norm Archer
DeGroote School of Business
McMaster University

Abstract

Chronic conditions often plague people, and many require frequent care and medications throughout their lifetimes. Mobile computer applications and sensors can be used to monitor and detect unusual chronic health conditions, to assist in self management of diseases, and to notify patients and their providers of trends and events in their health measures. This paper describes several applications of mobile systems that use the evolving role of Smart electronic Personal Health Record systems (SePHRs) to support these interactions. Some research challenges in this field will also be presented.

1 Introduction

Chronic conditions such as non-communicable diseases (eg., diabetes, cardiovascular disease and chronic obstructive pulmonary disease), long term mental disorders and certain communicable diseases such as HIV/AIDS are the leading causes of death and disability in the world. The World Health Organization reports that currently chronic diseases are responsible for 60% of the global disease burden. In Canada, the current cost of illness, disability and death due to chronic disease is over $80 billion annually. In response to this growing burden many countries around the world are taking a great interest in improving the management of chronic conditions [1]. Across Canada, every province and territory has programs for managing chronic disease and for training the chronically ill in self management. With this well-deserved attention to disease self management that will help reduce the overall societal burden of chronic disease, there are many ways in which technology could assist with the management process. In almost all such applications, it is important for the patient to have an associated repository where monitoring and self-entered event data can be stored, reviewed, and analyzed to detect trends and outliers. Electronic personal health record systems (ePHRs) can be used to meet this need. Patients have already begun to adopt them to keep track of health-related information and events such as prescriptions, doctor appointments, illnesses, etc. Some physicians also allow patient downloads of certain clinical records such as lab test results, etc., stored on physician electronic medical records systems (EMRs).

Consumers are generally enthusiastic about having access to their own electronic medical records. In a US online survey in 2007, 91% of respondents agreed that "patients should have access to their own electronic medical record," and 60% agreed that "the benefits of electronic medical records outweigh the privacy risks" [13]. Physicians are also in favour of their patients having access to their own records, but with some reservations [3]. Many patients use the Internet to e-mail their physicians, receive test results, schedule appointments, etc. But due to insufficient technical knowledge, most physicians and patients have an incomplete understanding of the capabilities of electronic health systems and the system interoperability needed to deliver needed information. A major drawback to existing ePHRs was found in an environmental scan of the ePHR field by Archer et al [4]. One of their major conclusions was that most of the ePHR systems that have been developed and studied are physician-oriented. Thus, many have not included additional patient-oriented functionalities that are needed for maximum effectiveness. Until such ePHR systems are provided for patients, it is unlikely that tangible and/or intangible improvements in health outcomes from patient health self-management will be demonstrated.

The purpose of this brief paper is to describe innovative extensions of ePHRs to develop Smart ePHRs (SePHRs), patient centric smart systems and repositories for managing data from patient and physician inputs, combined with monitored data from mobile devices from indicators such as heart rate, blood pressure, blood glucose, weight, and other measures relevant to patient health. Such systems, to be effective, must operate from mobile wireless platforms such as smart phones.

2 Smart Mobile Applications for Managing Health of Individuals

The Smart Internet [10] was recently introduced as the next Internet generation, to reflect the move of the Internet toward a more user centric network. Among other benefits, this concept will allow clients to select services that work best, based on client contexts and needs. To meet these goals, the underlying conceptual model and infrastructure of the Service Oriented Architecture (SOA) needs to be extended and modified. Emerging technologies such as Software as Services (SaS) and cloud computing propose to support applications as services which can then be distributed across the network and reused in other applications. This will result in an increase in the number of web services that offer similar functionalities. Discovering relevant services (service discovery) and choosing the best ones to meet client needs (service selection) will be more challenging. The following discusses a few related eHealth applications.

2.1 Virtual Remote Nursing

Traditional approaches to service selection comparison of web services are based on their descriptions published in service registries, including QoS and price/performance ratios, as well as adaptability. This information is generated by the service developer and may not be fully trusted by the client. Moreover, alternative services may exhibit different performance in different client contexts that cannot be determined accurately by service descriptors. In a recent paper Najafi et al [9], propose a novel service selection approach that compares alternative services based on their performance in a specific client context. For this purpose, the SOA reference model is extended with the addition of a "competition desk" component that manages competition among alternative services available to the service client. This assists clients in choosing services that best suits their needs. A generalized prototype system was developed for the proposed extended SOA model which supports two types of web services: traditional data services (where the service processes client data and resources completely at the server site) and the task services that are proposed in the paper (web services with the capability of processing client data and resources partially or completely at the client site).

A recent paper by Najafi et al [8] applies the novel service representative approach described above to a virtual remote nursing (VRN) system, with a virtual nurse agent installed on a patient's personal computer or smart phone to help manage the patient's health condition continuously. A variety of tasks can be assigned to the virtual nurse using a generic task definition mechanism. Here, a task is defined as a combination of medical workflow and operational guidelines such as those derived from evidence based medicine to guide physicians in the diagnosis of patients and the development of treatment plans [2]. Practitioners select the tasks that will help the patient follow particular treatment plans, and to indicate the patient's health status or to give the patient appropriate warnings, recommendations, and reminders in pre-defined situations. Combined with an ePHR that records monitoring and other data from the patient's electronic personal health record, this supports the use of the VRN by the patient and the health provider.

A prototype VRN is being developed, with plans to extend it to an SePHR operating on a smart phone platform. The smart version would gather data through both patient and environment monitors, along with patient-entered data, for either healthy patients or patients with a variety of diseases or conditions. Data gathered in this manner would be mined in the sense that correlations would be evaluated between context measures (such as activity and location) and patient measures (in the case of people with heart disease, heart rate and blood pressure) to determine the patient's functional response to specific contexts. Knowledge derived from this type of analysis would be helpful in determining limits [2] to patient exposure, as these would show up in continuing monitoring activities. Warning messages could be generated as the monitored results reach or exceed preset bounds. Data gathered in this manner would be very helpful to physicians in devising treatments for their patients, and it may be useful in other applications such as sports medicine for people wishing to keep fit [11]. Note that active monitoring of patients through devices such as "Holter monitors" to record heart electrical activities is well known. Linking these measures to context characteristics (e.g. the activity in which the patient is engaged) opens up new areas of investigation, such as the work by Mouttham et al [7] and Shahriyar et al [12] on personal health monitoring. Our work proposes to monitor environment (as context) and patient measures and then correlate them analytically to develop knowledge that will help both patients and physicians to recognize unusual activity and events and to take the necessary steps to remedy problems that may materialize as more information is recorded and analyzed. The system could be individualized to each patient through the knowledge gained from measures of patient response to day to day activities, thus responding to symptoms or measures that are unique to the patient.

2.2 Automated Message Triage

Tavasoli [14] developed an automated message triage (AMT) technique for supervised text mining to classify messages or short documents according to the semantic significance of their content. This technique uses a combination of several algorithms that are known to work well for classifying documents that may have a significant overlap in their content. Environments that display this type of overlap include e-mails, text messages, user opinions, requests for user support, etc. There are many possible applications of AMT in healthcare, and this high performance classifier has been demonstrated with messages exchanged between patients and physicians in a trial of a system that used an ePHR to gather data from patients with hypertension. The intent in cases like this is to provide 24/7 monitoring of patient messages and to redirect them according to their level of priority. The highest priority messages can be directed to the attending physician or emergency service, medium priority messages can be handled by nurses or other staff, and lowest priority messages may result in automated responses with educational or other online materials. The system, once trained on known priority messages, can improve its classification ability over time as new messages that arrive are reviewed by an expert to define an accurate classification for the AMT system.

One potential application of AMT is as an augmentation of a recently introduced system called *FrontlineSMS: Medic*, designed specifically with open source software that runs on a laptop [5]. It basically turns a laptop and a mobile phone into a central communications hub. Once installed, the program enables users to send and receive text messages to individuals or groups of people through mobile phones. This system was designed for, and is used in, developing countries to more efficiently help healthcare providers provide services to people within range of cell towers. When combined with ePHR capability and the AMT classifier, this system could become an SePHR that would reduce the load on physicians by distributing messages to others who can handle the specific requests or by automatically responding with information relevant to low priority requests. The associated ePHR database could also provide a source of continuity for online patient records, for patients and physicians alike.

The preceding advances in smart self managed healthcare with SePHRs represent some of the research we are involved in. With these working prototypes, more effort will follow in testing and validating real applications, making adjustments as needed to suit clients and their physicians, and ensuring reliability of the systems and validity of the results. Future research will also address issues related to homecare nursing applications.

3 Research Challenges

We have identified the following research challenges to the further development of SePHR applications such as those discussed above.

3.1 Establishing Effectiveness

It is a challenge for home healthcare workers to support chronically ill and frail people cost effectively at home, while government agencies are promoting homecare as a way to help reduce healthcare costs. Technological advances have generated new devices that may provide some of that support. Homes equipped with mobile applications, including sensors that monitor activities and send information to healthcare providers are called smart homes. However, many of these new technologies are used without appropriate evaluation. A recent review of the smart home literature [66] found a large volume of literature on the use of smart technologies in home healthcare, many of which measure user satisfaction, but none that tested their effectiveness. The challenge in this research is similar to problems faced elsewhere in the mobile health field, which is to establish the effectiveness of these interventions.

3.2 Privacy and Security

Maintaining security in mobile self management systems is a major challenge, but advances continue to occur in smart phone technology, making encryption of data and messages feasible. This requires further attention to system design, with a level of security that will maintain the desired level of privacy for users. Specific types of techniques for secure access to databases (eg., role-based, team-based, attribute-based) may be suitable for mobile self management systems.

3.3 Usability and Adoption

Although usability is often an afterthought, with SePHRs it must be kept in mind from initial concept development, and built into the design with a testing regime that accommodates the eventual users of the system. This means that a simple user interface is the first and foremost design principle. Otherwise adoption will not succeed.

3.4 Change Management

Even with the most usable interfaces, the organization that interacts with systems such as those described, including physicians, nurses, administrators, and others, may need to change workflows and procedures to accommodate these innovations. Moreover, stakeholders have to be convinced that such systems will benefit them as well as their patients. Accommodating to such changes must be considered and planned in advance, with the main consideration being to involve all the stakeholders in the selection and adoption process from the beginning.

Acknowledgments

The author would like to thank Dr. Kamran Sartipi, Ms. Shima Aghtar, Mr. Mehran Najafi, and Mr. Amir Tavasoli for their work on the research reported in this paper. This research was supported by a grant from the Natural Sciences and Engineering Council of Canada.

About the Author

Dr. Norm Archer is a professor emeritus in the Information Systems Area of the DeGroote School of Business, and a special advisor to the McMaster eBusiness Research Centre (MeRC). His main research interests are in eHealth, and he was a co-founder of the MSc eHealth program at McMaster University, a collaborative initiative of the DeGroote School of Business, the Faculty of Health Sciences, and the Faculty of Engineering. It is now in its third year of operation.

References

1. Alberta Health Services. *Chronic disease management overview http://www.calgaryhealthregion.ca/programs/cdm/ (accessed Oct. 7 2010)*. 2010.
2. Anonymous, *Hypertension guidelines NGC-5081*, in *Guideline Summary*. 2005, Kaiser Permanente Care Management Institute: Oakland CA. p. 9.
3. Archer, N. and U. Fevrier-Thomas, *An empirical study of Canadian consumer and physician perceptions of electronic personal health records*, in *ASAC 2010*. 2010, Administrative Science Association of Canada: Regina, SK.
4. Archer, N., et al., *Electronic personal health records: An environmental scan*, in *McMaster eBusiness Research Centre Working Paper #34*. 2010, DeGroote School of Business, McMaster University: Hamilton, Canada.
5. Banks, K. (2009) *FrontlineSMS: About the project http://www.frontlinesms.com/abouttheproject/about-the-project/ (accessed October 6 2010)*.
6. Martin, S., et al., *Smart home technologies for health and social care support.* Cochrane Database of Scientific Reviews, 2010(9).
7. Mouttham, A., et al., *Event-driven data integration for personal health monitoring.* Journal of Emerging Technologies in Web Intelligence, 2009. **1**(2): p. 110-118.
8. Najafi, M., et al. *Virtual remote nursing system* in *First IEEE International Workshop on Consumer Health Platforms, Services, and Applications*. 2010. Las Vegas, Nevada: IEEE.
9. Najafi, M., K. Sartipi, and N.P. Archer, *Web service competition: A service selection approach*, in *Submitted for Presentation*. 2010.
10. Ng, J., M. Chignell, and J. Cordy. *The smart Internet: Transforming the web for the user*. in *CASCON '09: Proceedings of the 2009 Conference of the Center for Advanced Studies on Collaborative Research*. 2009. New York, NY: ACM.
11. ORTO. *Systems for body-state diagnosis http://www.cardiostep.com/en/research.shtml (accessed Oct. 7, 2010)*. 2010.

12. Shahriyar, R., et al., *Intelligent mobile health monitoring system (IMHMS).* International Journal of Control and Automation, 2009. **2**(3): p. 13-27.
13. Steinbrook, R., *Personally controlled online health data: The next big thing in medical care?* New England Journal of Medicine, 2008. **358**(16): p. 1653-1656.
14. Tavasoli, A., *Automated Message Triage: A Proposal for Supervised Semantic Classification of Messages*, in *Computing And Software Department*. 2010, McMaster: Hamilton, Canada. p. 100.

Smart Health

&

Wellbeing

Yelena Yesha

University of Maryland

Baltimore County

Towards

Patient-Centric Healthcare

Igor Jurisica

Ontario Cancer Institute

Sociality

&

Privacy

Persons, Communities and Artifacts on The Web

Eleni Stroulia

Department of Computing Science,
University of Alberta, 221 Athabasca Hall,
T6G 2E8, Edmonton AB, Canada
stroulia@ualberta.ca

Abstract. In this paper, we review some recent trends in the evolution of social systems on the web today and we discuss how these systems may be brought to bear to improve existing services and to enable new innovative services for web users.

Keywords: Groupwork, Groupware, Social networks, Web 2.0, Web 3D

1 Background

The World Wide Web was conceived and born out of the desire to support information exchange, communication and collaboration. In its 30-year history (and it is flabbergasting to think about how short, in terms of time, this history is and how dense, in terms of events and innovations) it has more than fulfilled its promise and vision while at the same time undergoing three interesting transformations.

In the beginning, the objective of the web community was to enable document publishing and to advance large-scale information communication. The first beneficiaries of this platform were the academic and research community who had the knowledge and skills (a) to develop "web portals" even without any development tools and (b) to access the published information through the original crude client applications. Through this activity, the first broadly usable clients and web-development toolkits were developed and gave rise to portals supported by traditional and new content owners, such as mainstream print publishers (MIT's Tech newspaper in 1991, BBC's TV program in 1994, and the Clinton White House in 1994) and new content providers (Yahoo in 1994). In this stage, the web was ***a web of information*** broadcasted by few to many.

The second phase transition in the Web's history was brought by the advent of ecommerce sites (Amazon and eBay in 1995), which gave rise to ***the web of applications;*** the web became a ubiquitous platform through which to deliver innovative services. The number of providers increased dramatically as the community became ever more creative about the types of services that could migrate to the web. The number of consumers also exploded with the increased availability of user-friendly browsers, search engines (Alta Vista, the first multilingual engine, was launched in 1995) and email-service providers for individuals (Hotmail was launched in 1996). Still, the communication model was broadcasting by relatively few to many.

This changed with the advent of bulletin boards, originally associated with ecommerce web sites, and wikis and blogs, easy to use publication tools for individuals. These tools brought about ***the personal web***, a continuously available whiteboard, hosting everyman's opinions and personal expressions, across the world.

And as the tools for searching, tagging, visualizing and connecting personal posts, published through any of the multitude of available platforms, became increasingly available, ***the social web*** emerged. Today each one of us is linked to a multitude of others through our on-line presence: to the authors of the blogs to which we comment, to the other buyers of the products and services we have bought, to the members of our professional communities (linked-in and ning), to the people whose micro-blog postings we follow (twitter), to our on-line friends (facebook), to the members of our virtual-world communities (second life), and to the users of the on-line tools we use.

Clearly, the original web vision, of supporting collaboration, has been evolving throughout these phase transitions, and today, it appears that the potential for innovative modes of web-enabled collaboration has reached new heights. It is in fact at the core of the "smart planet" ***interconnectedness*** vision, which includes (a) data, (b) system and (c) people interconnectedness.

In our work, motivated primarily by the need to support collaborative software development, we have developed a family of systems for supporting, managing and analyzing different types of collaborative activities. In the rest of this paper section, we review this family of systems and we place it in the context of related work (Section 2). Next, we identify what we believe are some interesting questions in terms of which to understand and analyze social systems (Section 3) and we review our work on SociQL, a social query language designed to support the expression of such analyses (Section 4). Finally, we discuss some ways in which social systems can be brought to bear in service delivery (Section 5) and we conclude with some thoughts on what we expect to be the next important innovations to come (Section 6).

2 Collaboration in the Social Web

In the past several years, our team has developed four different web-based systems to support, manage and/or analyze four different types of collaborative work. Looking back through this work, we have attempted to place it within a coherent conceptual framework by categorizing each tool in terms of two dimensions: (a) the type (and flexibility) of collaborative practices they support and (b) the type of technology/platform they assume.

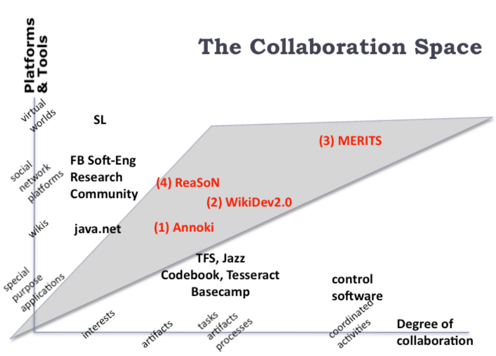

Figure 1: Collaboration Tools in a Two-Dimensional Space

As shown in the vertical axis of Figure 1, the adopted platforms range from task-specific to general-purpose, with the latter category including wikis (and blogs), social-network platforms and virtual worlds. In the horizontal axis of Figure 1, we have identified several interesting spots in the continuum of collaborative practices, from simply establishing communities with common interests, to groups of people sharing artifacts of interest, to teams that collaborate by exchanging artifacts and information according to established process, to very regimented workflow tools that enact well-defined processes to which people contribute well-structured artifacts.

2.1 Annoki [1,8]

Annoki was built on top of the popular MediaWiki, to support the collaboration of our research team. We chose MediaWiki as the platform because its default features fulfill many of our original requirements for our envisioned research-collaboration tool. First, it provides "user" pages, for the personal use of the wiki members, and "regular" wiki pages where content is collaboratively edited. Second, it has a "discussion" page for each "regular" page, thus enabling a distinction between "content" and "reviewer's comments" among the collaborators. Third, it supports concurrent editing of pages (with the multiple versions getting merged a-la SVN) and notifications of users when a page of interest changes; these two features enable a tighter, more synchronous coordination among collaborators. Finally, it supports templates, so that in addition to free-formatted pages, structured information can be collected.

To the MediaWiki features, the Annoki toolkit adds the following set of extensions.

Namespace-based access control: Each Annoki user has an associated namespace and all pages he creates belong in this namespace. Group namespaces can also be defined to organize wiki pages that "belong" to a group of users. Pages belonging to a "public" namespace are visible to all. In this manner, layers of protection can be supported for personal, project-specific, organization-related, and publicly accessible content.

Annotations and visual editing of template instances: To enable lightweight cross-referencing of pages, users can annotate pages with their own tags. In this manner, users can superimpose a personal layer of their own on the wiki resources. Users can also create and edit pages based on templates using a graphical editor, removing the burden of writing wiki page code.

Visualizations: Annoki is equipped with two types of rich, interactive, Ajax-based visualizations: WikiMap and wiEGO. WikiMap is a visualization of the whole wiki structure (users, pages, links among pages and authorship relations between users and pages). The set of wiEGOs are visualizations of the semantic structures implicit in a set of special template-based pages corresponding to concepts in Blooms taxonomy (i.e., tree, topic, persuasion, brainstorm, story, and decision maps, as well as flowcharts).

Collaboration and contribution analysis: extending the default differencing capability of MediaWiki, Annoki supports analysis of the page edit history at the level of sentences, and collects metrics of each user's contribution to each page and to the wiki as a whole in terms of sentences added, deleted, and edited.

2.2 WikiDev2.0 [3,4,6]

The WikiDev2.0 tool for collaborative software development was conceived as a lightweight platform through which to integrate information about various software artifacts produced in the variety of tools used by the software team (code, documentation, communication messages etc), to analyze this information in order to infer interesting relations among these artifacts, the team members and their activities, and to present views on this information that cut across the individual tool boundaries.

The ***code and communication clustering*** process of WikiDev2.0 consists of the following steps. The first step involves parsing of all the textual information associated with the input information feeds, to recognize mentions of team members (their names, nicknames, or IDs) and software artifacts (classes, methods and interfaces). The recognized references introduce the explicit relations between people, code and communication artifacts. A subsequent step calculates the implicit relations based on triangular inequality thus providing further insights about hidden dependencies among these artifacts. Using this information, one can see who works on what artifact currently, who has discussed a specific artifact that should be potentially consulted about changes to it, and how a member's own work might affect other people's work.

The ***syntactic-semantic text-analysis feature*** of WikiDev2.0 is meant to further enhance the ability of the tool to recognize relations among people, code and communication artifacts, implicit in the large amounts of textual information collected through the software process. The process consists of (a) a syntactic parsing stage for all textual content in WikiDev2.0, (b) an annotation stage, where the syntax trees of the parsed sentences are annotated with semantic information (such as team-members' names and code artifact names), and (c) a pattern-matching query stage that extracts subject-predicate-object triples from the annotated parse trees, corresponding to relations such as "who worked on what", "who has experience in what" etc.

We have developed ***a variety of visualizations*** in WikiDev2.0 to communicate the state of the project to team members and managers. Traditional line- and graph-based charts communicate the amount, type and frequency of team-members' activities. The UMLViewer resents an UML-like view of the code artifacts, annotated with information about their developers and evolution.

Finally, ***WikiDev2.03D*** is an extension built in the Open Wonderland virtual world, which visualizes the discovered clusters adopting a 3D city metaphor. This virtual-world view of the project can be visited and discussed by multiple interested parties at the same time, thus enabling a shared understanding of the software project.

2.3 MERITS [2,5]

The MERITS system is the third of our collaborative-work support tools and it focuses on activities that are more complex and involve, in addition to information exchange, interactions among people and between people and the real world. To that end, it combines the immersive, collaborative potential of virtual worlds with BPEL-based process specification to enable (a) instructors to specify educational scenarios, and (b) students to experience those scenarios in a realistic, interactive manner.

The MERITS framework offers two important features. The first involves a method and tool support for ***specifying complex collaborative processes,*** including tools for specifying the behavioral capabilities of the various roles in the process in terms of web services invoked by avatar actions in the virtual world, as well as developing behavioral scripts for real-object simulacra in the virtual world.

The second important feature is a ***comprehensive action-recording*** tool that produces a compact trace and a synchronized trace of all in-world actions, which can then be parsed to identify interesting action patterns.

2.4 ReaSoN [7]

Shifting focus from supporting to analyzing collaborative activities, we developed ReaSoN (again based on Annoki), a comprehensive set of tools for visualizing and exploring the social networks, implicit in academic research practices. In doing so, ReaSoN contributes to the understanding as well as fostering of the social networks underlying academic research.

In terms of ***visualizations***, ReaSoN offers specially structured pages to communicate information about individual and collections of *publications*, *authors*, the *communities* around conferences and journals, the *keywords* of publications, and the geographical distribution of people, keywords and communities.

More interestingly, ReaSoN also provides infrastructure for asking customized queries about researchers, their collaborators, their publications and their citations. The results of a query can be visualized in tabular form, explored in the WikiMap graph format (making visually explicit the network of authors and publications) or plotted on a map (visualizing the geographical relations among people organizations and their research activities).

3 Analyzing the Social Web

Developing and reflecting upon the four systems we described in the last section, we have come up with a set of research questions (and associated technical challenges) that cut across most (all?) web-based social systems today. We review these questions in the remainder of this section, organized in two different groupings of "analysis questions" around social systems and possible "services supported" by social systems.

Today, there exists a plethora of social-networking sites, each one supporting different means of "connecting" among members and catering to different demographics. Some sites enable bi-directional connections, like Facebook, where others enable directed connections, like Twitter. MySpace caters to a younger demographic than Facebook, which in turn is surpassed in popularity by Orkut in Brazil. In addition to these "superficial" differences, each of these social networks encourage different types of communications. Facebook appeals to people who want to keep in touch with family and friends where twitter seems to be the medium of choice for people to share and access information from a wide variety of channels. Facebook favors deeper connections and enables the organization of these connections in groups so that different personas can be projected to each of them. Twitter, on the other hand, encourages maximization of connections (followers) and enforces a single persona its users who cannot distinguish their followers in groups. Clearly these differences deserve deeper analysis; in the mean time, all of these networks share three important concepts, i.e., community, contribution and influence.

3.1 Recognizing Communities

Groups of collaborating people are not uniformly cohesive. Some members are more highly connected to each other than to the rest of the group. This is a corollary to the "homophily" phenomenon. Homophily is the tendency of individuals to associate and bond with similar others. Individuals in homophilic relationships share common characteristics (beliefs, values, behaviors, etc.) that make communication and relationship formation easier. In principle, graph algorithms for connected-components' recognition can be applied to recognize such "cliques". Alternatively, domain-specific notions of subcommunities can be defined.

Let us review the issue of "recognizing communities" in the context of our systems above. In ReaSoN, we have analyzed the communities of authors who have published in specific conferences over a period of time and the intersections of these communities with each other. In WikiDev2.0, we have analyzed the email communications of team members to recognize subgroups who have communicated most frequently with each other. We have also clustered communication artifacts around the code artifacts they relate to, and by implication the authors of these code and communication artifacts. In Annoki members belong in project-related Namespaces, which essentially define the communities of members and documents that are associated with a project; thus there seems to be no point in recognizing implicit subcommunities. In MERITS workflow-defined simulations, the activities of the various participants are understood in terms of the workflows they enact; however, in cases of more open-ended activities, special-purpose relations (like communication) could be defined in terms of which to recognize dense subcommunities.

3.2 Recognizing Contribution

As the collaborating community increases, the roles of individuals become blurred and unclear. In WikiDev2.0, for example, most teams consist of four to six developers (plus TAs and instructors). Contrasting this to the about 200 members of the Annoki installation for the software-engineering group at the University of Alberta, it becomes clear that the latter community is much more complex than the former. Recognizing the contribution of individuals in the latter context becomes challenging.

MediaWiki, as well as most wikis, offers a differencing capability, which summarizes the contribution of an individual to a specific version. Annoki provides a more sophisticated contribution analysis and visualization tool, which summarizes the contribution of an individual to a wiki page over its lifecycle. WikiDev2.0 implicitly recognizes contribution in terms of frequency of SVN commits, wiki-page edits, and email communications. ReaSoN offers a variety of bibliometrics-inspired statistics to measure the "importance" of each author, including their h-index and various pagerank calculations of the influence of their papers to other papers and their corresponding authors through citations. MERITS does not offer an explicit contribution-measurement solution since it is a framework, and contribution measures, in general, have to be aware of the nature of the collaboration activity. Instead, through its recorded activity logs one can define metrics of interest based on the participants' in-world activities and measure contribution in different ways. For example, one can imagine that it would be interesting to identify the persons who talked the most during a session or the person who made the most interactive gestures (like shaking hands for example) with others.

This discussion assumes that "importance" is semantically equal to "contribution" which is not necessarily the case. Domain-specific importance metrics can be based on different person attributes, but contribution appears to be a cross-domain importance metric.

3.3 Recognizing Influence

Related to the concept of contribution is the concept of influence. Within a collaborating community, people influence their collaborators through their contributions. Not all contributions however are equally likely to be consumed and to influence other people's contributions.

In ReaSoN, we measure influence through pagerank calculations over the implicit coauthorship and citation networks. Through these metrics, one can recognize authors with broad co-authorship networks, i.e., who have written papers with many other authors who have similarly written papers with many others etc., as well as authors with broad citation networks, i.e., authors whose papers have been cited by many authors whose papers have been cited by many other papers etc. In WikiDev2.0, the clustering process implicitly attempts to recognize the members' influence to code artifacts by collecting references of other materials, associated with team members, to these artifacts. It does not offer however any insight on how to compute any type of transitive closure of these relations.

4 SociQL

Aiming at understanding the various types of collaborative work exemplified by the above systems and at supporting a general conceptual framework in which to address the above research questions, we are now working to design asocial query language. SociQL is a query language, and an associated prototype implementation, that supports for the representation, querying and exploration of disparate social networks.

Unlike generic web query languages, SociQL is designed to support the examination of such sociological questions, incorporating social theory and integration of networks that form a single unified source of information. In sociology, object-centered sociality characterizes social relations between individuals by means of objects. In this setting, specific constitute evidence of social relations (Knorr-Cetina 1997). Essentially, while recognizing the social interaction between *individuals*, this theory exalts the role of *specific objects* as the reasons why social actors affiliate with each other.

For this reason, we define SociQL's data model around the concept of an **object**. For instance, in the context of ReaSoN, we have that a paper (an object) connects the researchers who authored it; similarly, a publication venue (an object) connects authors who publish their work in it. In the MediaWiki-based, Annoki and WikiDev, the wiki pages are the objects that connect the pages authors. In MERITS, the avatars are connected through the simulacra of the real-world objects they manipulate as well as through their communication objects, i.e., their text and voice utterances.

In our model, both objects and relationships are described by **properties** (actual data), such as the name of an author or the date in which a citation is made from a paper into another. We also distinguish the **context** in which properties are defined to describe the objects and relationships. For instance, the same query might return different email addresses for the same individual depending on the context in which the query is asked (professional or personal). In practice, each context will correspond to different social network system—thus, each context may have its unique data access methods and privacy restrictions, which complicates query processing to a great extent (as discussed below).

As it turns out, this problem is extremely hard to solve in practice. In order to correctly interlink the different communities, different social network sites describing the same object would have to refer to it with a globally consistent identifier. In practice, however, each site has its own local identifier, unique only in its particular context. This practice results in a proliferation of identifiers that make it harder to merge social networks.

5 Services with Social Support

As the number and types of social systems increase and so is their membership, the question becomes to identify the means through which they can be brought to bear in delivering novel and/or improved services.

An interesting new technology than can provide a catalyst for the deployment of social-network information to consumers is the combination of QR tags and the availability of tag readers on almost all new mobile devices. By scanning the QR tags annotating real products and business cards, mobile apps can inform the individuals' networks of their real-world consuming behavior and social interactions. In this manner, the social network itself can seamlessly expand through traditional real-world practices (like business-cards' exchanges) and the word-of-mouth advertisements for products and services can efficiently travel through it. Similarly, information about individuals' entertainment choices can be propagated through their networks as, increasingly, we are consuming entertainment, games, audio and video, through the Internet. As more information is shared by the network, collaborative filtering becomes more effective in advising the network members about what their connections buy, play, listen and watch. And to the extent that more network members choose to make similar choices they can negotiate better prices and improved quality for their "group buying".

6 Summary

In this paper, we discussed our recent work on four collaborative/social systems, and our more recent work on a social query language designed to express the types of questions that users (and applications) may want to answer in the context of such systems. Further, we reviewed the types of analyses that we believe are relevant in the context of social systems and the ways in which social systems can be deployed to improve current services to Internet users and to enable new innovative services. This is clearly an active and fascinating area with a huge number of open questions and substantial opportunities for the development of innovative intelligent services.

Acknowledgements

This work was supported by NSERC, iCORE, IBM, and was conducted in collaboration with B. Tansey, K. Bauer, M. Fokaefs, D. Chodos, D. Serrano and D. Barbosa.

References

1. O. Arazy, Stroulia E., Ruecker S., Arias C., Fiorentino C., Ganev V., and Yau T., Recognizing Contributions in Wikis: Authorship Categories, Algorithms, and Visualizations, Journal of the American Society for Information Science and Technology (JASIST), 61(6):1166-1179, Mar. 2010.
2. D. Chodos, P. Naeimi, E. Stroulia: A simulation-based training framework for health-science education on video and in a virtual world, Journal of Virtual Worlds Research 2(1), Apr. 2009.
3. M. Hasan, E. Stroulia, D. Barbosa, M. Alalfi: Analyzing Natural-Language Artifacts of the Software Process, ICSM 2010, ERA track, 26th IEEE International Conference on Software Maintenance, September 12-18, 2010, Timişoara, Romania.
4. M. Fokaefs, D. Serrano, B. Tansey and E. Stroulia: 2D and 3D Visualizations in WikiDev2.0, ICSM 2010, ERA track, 26th IEEE International Conference on Software Maintenance, September 12-18, 2010, Timişoara, Romania.
5. D. Chodos, E. Stroulia, P. Kuras, M. Carbonaro, and S. King: MERITS Training System - Using Virtual Worlds for Simulation-based Training, CSEDU 2010 (the International Conference on Computer Supported Education), April 6-9 2010, Valencia Spain.
6. 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.
7. V. Ganev, Z. Guo, D. Serrano, B. Tansey, D. Barbosa, E. Stroulia: An Environment for Building, Exploring and Querying Academic Social Networks. The International ACM Conference on Management of Emergent Digital EcoSystems (MEDES) 2009, October 27-30, Lyons, France.
8. B. Tansey, E. Stroulia: Annoki: A MediaWiki-based Collaboration Platform, In. Web2SE: First Workshop on Web 2.0 for Software Engineering, ICSE 2010.

Paying Attention in Meetings: Multitasking in Virtual Worlds

Kelly Lyons, University of Toronto

Henry Kim, York University

Saggi Nevo, SUNY Albany

ABSTRACT

In this paper, we present the results of interviews with 23 corporate virtual world users which show how the use of virtual worlds for distributed meetings may help to overcome some of the problems associated with other kinds of technologies used to support meetings over distance. In particular, the results of our study suggest that the extent of multitasking enacted during virtual worlds meetings is in between face-to-face meetings and teleconferencing. When participants multitask, they simultaneously engage in other activities (e.g., e-mailing and texting) while attending the meeting, thereby paying less attention to the meeting's agenda. It has been argued that multitasking can be beneficial in some cases yet detrimental in other cases. Technologies that support distributed meetings should, therefore, afford an opportunity to engage in some degree of multitasking but not so much as to distract them from the meeting goals. Based on a series of interviews with 23 users of corporate virtual worlds, our study suggests that virtual worlds are capable of providing this desired balance.

KEYWORDS: Virtual Worlds, meetings, collaboration, multitasking.

INDEX TERMS: H.4.3 [Information Systems and Applications]: Communications Applications -- Computer conferencing, teleconferencing, and videoconferencing

1 INTRODUCTION

With the increase in globalization, an increasing number of companies conduct business across distance. More companies than ever are distributed across geographic boundaries and time zones, and many employees are working from home or, otherwise, telecommute. Faced with adverse economic conditions, companies are attempting to save money by reducing the number of meetings involving travel and overnight stays and are further motivated by a desire to reduce their carbon footprints.

This means that meetings are increasingly taking place over distance, supported by some combination of technologies including teleconferences, video conferences, electronic meeting software and, more recently, virtual worlds (e.g., Kahai, Carroll, & Jestice, 2007; Lucia, Francese, Passero, & Tortora, 2008).

In this paper, we look at the organizational use of virtual worlds for meetings and focus on the extent of multitasking exhibited by the individuals attending the meeting. Multitasking has been associated in the past with meetings over distance supported by technology. Mark, Grudin, and Poltrock (1999) report on the use of desktop conferencing with application sharing to support meetings over distance by four teams at a major corporation. Among their findings, they describe that many people reported multitasking during the meetings and most of the people in their study found multitasking to be advantageous in that they could accomplish other work and attend more meetings. However, some people in their study reported that multitasking was a "distraction and detriment" (Mark, Grudin, & Poltrock, 1999). Tang (2005) discusses how to design technology to reduce social awkwardness associated with multitasking computer interactions with other activities such as participating in meetings.

We characterize multitasking in meetings as not paying full attention to the meeting's proceedings. It can mean attending to email or instant messages during meetings, reading unrelated or even related material during a meeting, or engaging in any other activity that is not part of the discussion or current agenda of the meeting. People multitask for a number of reasons such as productivity (Mark, Grudin, & Poltrock, 1999) or not wanting to miss anything (Stone, 2006). Researchers have noted that multitasking has become so common that it virtually a social norm (e.g., Tang, 2005).

Our study shows that multitasking in meetings takes place more often in teleconferences than in face-to-face situations. This may be because people feel a level of social awkwardness when multitasking in a face-to-face meeting (Tang, 2005) or may be due to social concepts of awareness and accountability (Erickson & Kellogg, 2000). An interesting question is whether multitasking has a positive or negative impact on various aspects of meetings (for example, outcomes, effectiveness, efficiency, or personal relationships of participants). On the one hand, multitasking can increase effectiveness by allowing participants to pull in needed information or engage with someone not currently in the meeting in order to move an agenda forward or to fill in gaps in knowledge. On the other hand, multitasking can cause people to be distracted and not give sufficient attention to the meeting and the issues at hand. There is also a social component associated with multitasking that might cause others in the meeting to feel that their colleagues are not giving sufficient attention to the issues on the agenda.

To further understand the advantages and disadvantages of multitasking in distributed meetings and to understand the impact of the enabling technologies, we conducted semi-structured interviews with virtual worlds' users at a large multinational enterprise (LME). Several themes emerged from our analysis of the interviews. In particular, the use of virtual worlds for meetings emerged as a prominent theme. The identification of this type of usage of virtual worlds promoted a second round of interviews in which we asked interviewees specific questions about business meetings in general, meetings in virtual worlds in particular, and multitasking in meetings.

The results of our interviews further substantiate prior research (Mark, Grudin, & Poltrock, 1999; Stone, 2006; Tang, 2005) which found multitasking to be: (1) a common occurrence in technology-supported meetings over distance and (2) considered both a positive and a negative activity. Our study of meetings in virtual worlds shows that virtual worlds provide an environment in which individuals' levels of multitasking, while not quite the same as that in face-to-face meetings, are greater than in teleconferencing.

The rest of the paper is organized as follows. We present a review of literature of virtual worlds, technology-supported distributed meetings, and multitasking. We then present our research methodology followed by our findings about meetings in virtual worlds and multitasking in meetings. This is followed by a brief discussion of the key findings. We then conclude with suggestions for future research.

2 LITERATURE REVIEW

The research on virtual worlds has been undertaken by different disciplines with a focus on a variety of issues. Early research focusing on building and improving virtual 3D environments came out of computing science, engineering, and other applied technological sciences (eg, Mackinlay, Card, & Robertson, 1990). More recently, researchers have begun to examine the use of these 3D virtual environments in a variety of different contexts. Researchers in sociology and psychology are primarily interested in understanding how social behavior norms in virtual worlds differ from the real world (Yee, Bailenson, Urbanek, Chang, & Merget, 2007). There is an interest in political and legal issues such as liability, regulation, intellectual property, and economic and monetary systems in virtual worlds (Balkin & Noveck, 2006; Mayer-Schonberger & Crowley, 2006; Sipress, 2006; Soraker, 2008). Researchers have studied virtual worlds from the perspective of market research (Hemp, 2006), advertising (Barnes, 2007), eCommerce (Holzwarth, Janiszewski, & Neumann, 2006), organizational issues (Nebolsky, Yee, Petrushin, & Gershman, 2003), industrial engineering (Edwards, 2006), and management information systems and services (Goel & Mousavidin, 2007; Kadavasal, Dhara, Wu, & Krishnaswamy, 2007). There is also increasing interest in the use of virtual worlds for education (Boulos, Lee, & Wheeler, 2007; Bronack, Riedl, Tashner, & Greene, 2006). A detailed literature review (including applications of virtual worlds in business, education, law, and the social sciences) can be found in (Messinger, Stroulia, Lyons, Bone, Niu, Smirnov, & Perelgut, 2009).

Recently, the topic of meetings in virtual worlds has been identified as an important area for future study (Bessière, Ellis & Kellogg, 2009; Hendaoui, Limayem &Thompson, 2008; (Kahai, Carroll & Jestice, 2007). While there are relatively few studies about the use of virtual worlds for meetings over distance, there are numerous studies into other kinds of technology support for meetings over distance. Some work looks at meetings in general and the effect of distance on various aspects of meetings and collaboration (Olson & Olson, 2000; Bos, Shami, Olson, Cheshin, & Nan, 2004; Siina 2007). There has also been a large amount of research introducing novel tools and technology to support meetings over distance (Erickson & Kellogg, 2000; Lucia, Francese, Passero, & Tortora, 2008). Other work has focused on the use and effectiveness of existing technology for supporting meetings over distance (Cameron & Webster, 2005; Isaacs & Tang, 1994; Mark, Grudin, & Poltrock, 1999).

Multitasking has been studied in a variety of ways such as in general work environments (González & Mark, 2005), as a result of ubiquitous computing (Tang, 2005), and in meetings (Mark, Grudin, & Poltrock, 1999). To the best of our knowledge, no study, to date, has investigated multitasking in meetings conducted in virtual worlds. In this paper, we look at the use of virtual worlds to support meetings over distance and, in particular, the effect of virtual worlds on the phenomenon of multitasking in meetings.

3 RESEARCH METHODOLOGY

The data collection was carried out in two phases: we first conducted open, semi-structured interviews with 23 corporate users of virtual worlds and analyzed the results of those interviews; we then asked follow-up questions specifically addressing the notion of meetings in virtual worlds and analyzed the answers to those questions.

3.1 Data Collection and Analysis -- Phase I

The semi-structured, face-to-face, and phone interviews with key informants were used to develop a better understanding of the emerging phenomenon of virtual worlds, seeking to gain insights on how these technologies are used to carry out organizational tasks.

To maximize the breadth and potential to generalize the results from this first phase of our study, we interviewed key informants who use virtual worlds for different reasons and who have varying degrees of experience with virtual worlds. Overall, 23 individuals from a large multinational enterprise (LME) were interviewed. Just under half (48%) have been working at LME for ten years or more and 22% have been with LME less than five years (the rest had between five and ten years experience at LME). Our sample was 52% male with 43% in a managerial position. The participants came from across the LME organization, representing a variety of roles and positions (e.g., consultants, directors, project managers, and designers).

Typically, the interview location was a conference room at LME, and when interviewees were teleworkers or lived in a different country, interviews were conducted by phone. The average interview lasted 45 minutes and interviewees were asked (and granted permission) to record the interview. The first five interviews were conducted by all the authors and the rest were divided between two of the authors who then conducted one-on-one interviews. This approach was taken to allow for faster data collection while reducing the risk of bias and ensuring the consistency of the interview.

After conducting 20 interviews we felt that a sufficient level of saturation in the data was reached (Strauss & Corbin, 1994). As a measure of percussion, we conducted three additional interviews and decided to stop the data collection process when it became apparent that no new issues were raised by the interviewees and sufficient convergence was obtained.

The interviews were transcribed by a professional transcriber and then analyzed and coded by the authors: the data were read, parsed, and categorized into categories. In most cases, multiple confirming comments from different key informants were used to determine the coding value. To establish a common data analysis approach, the authors agreed on a set of guidelines and then continued to code independently. Findings were then used to populate a single database.

We identified important variability among the interviewees: some of our informants were experts who not only use virtual worlds for internal business purposes but use them in work with clients as well, whereas others identified themselves as novices and noted that they use virtual worlds intermittently; however, each informant has participated in at least one corporate project on virtual worlds.

3.2 Data Collection and Analysis -- Phase II

During phase I, the use of virtual worlds for meetings was identified as a major theme. Accordingly, we returned to the participants from Phase I one year later to ask them some specific questions about meetings and technology used to support meetings. We contacted them by email reminding them of the previous interview, letting them know that the theme of meetings in virtual worlds had come up in our interviews, and asking them to answer several follow-up questions. The follow-up questions were presented on-line and automatically collected in a database. Some of the people we originally interviewed had left LME by the time phase II was conducted. To make up for the loss of informants we invited additional employees of LME to answer our questions. Overall, 19 people responded to the follow-up questions. Since the original group and the secondary group do not completely overlap in terms of the participants, it was deemed best to decouple the interview results of the two phases.

More often than not, knowledge workers within organizations are geographically distributed yet the need for them to meet, make decisions, get to know one another, and collaborate remains at least as important, if not more, than it did in the past (Johnson, Manyika, & Yee, 2005). At the same time, organizations are cutting travel for face-to-face meetings in order to save costs and reduce their carbon footprint. A recent Linden Lab case study about Second Life reports that IBM estimates savings of over $250,000 in travel and venue costs and additional savings of $150,000 in productivity gains (since the virtual meeting participants did not have to leave their computers to participate and return to work quickly) (Linden Lab, 2009).

Our interviewees confirmed that travel has been reduced at LME with seven informants describing fewer opportunities to travel face-to-face to meet with colleagues. One informant said, “You know, obviously this is probably an industry wide phenomenon but [people who work at LME] don’t tend to travel a lot just to have face-to-face business meetings. We’re really encouraged to use other means, whether it’s telecoms or videoconferencing,” with another saying “And also because we’re now a global business, we work ... virtually. We’re keeping down on travel costs.” Two informants mentioned the importance of not only cutting costs but also reducing carbon emissions, one saying, “And that’s important because you know, carbon footprints are important. Cost reduction is important for the enterprise,” the other saying, “the calculation was that for a given event ... that cost $180,000 ... that meeting was 56 metric tonnes per event of CO2.”

When we asked our 19 informants how often they participate in meetings face-to-face versus over distance, 17 out of 19 are often or always in meetings supported by teleconferences and nine informants rarely or never meet in person. None of the people we asked meet in person exclusively, and of the eight who said they meet in person often, all but one also meets by teleconference or with eMeeting technology often. The people in our study are often meeting with colleagues over distance using a mix of mediating technologies.

Not only are people meeting colleagues over distance but there is often a significant difference in time zones of the meeting participants. We found that 13 out of our 19 informants sometimes or often participate in meetings where the time zone difference is 4-6 hours. Six informants said they sometimes participate in meetings where the time zone difference is 7-9 hours and five people said that sometimes the time zone difference is greater than nine hours. Two people indicated that they often participate in meetings with a 7-9 hour time zone difference and one person said the time zone difference in their meetings is often greater than nine hours. Our group of informants frequently participates in meetings with participants from across the world.

There is a need to figure out how to support these meetings better. A suggested way to do so is by using virtual worlds (Kahai, Carroll, & Jestice, 2008; Bessière, Ellis, & Kellogg, 2009). Informants in our study indicated to us that meetings are promising business use of virtual worlds. Of the original 23 people we interviewed, 16 people brought up or introduced the idea of conducting meetings in virtual worlds. In all of these cases, the informants did so without any suggestion of meetings in virtual worlds by the interviewer. It was most often brought up by informants in response to the questions:

“What benefits do you think virtual worlds bring (or can bring) to organizations?”;

“How do you think virtual worlds can be used for work?”; or

“How do you use virtual worlds?”

In one case, an informant discussed setting up communities of practice in a virtual world as an example of their own use of virtual worlds and went on to describe meetings of this community held inside the virtual world. One informant talked about participating in meetings in virtual worlds when asked how they first got involved in virtual worlds and another described meetings in virtual worlds when asked to provide a narrative of how they spend their typical day. Another informant talked about how they started getting involved personally in virtual worlds a year prior to the company getting involved. When asked if the informant was still involved for personal use, the response was, “... I find that I have less time to do anything in there on a social or after work time period because I already am in there for meetings and other events that I do.” One individual described setting up business meetings in virtual worlds for clients when asked to describe their role in the company.

While more and more teams and work groups are geographically distributed with fewer opportunities to meet face-to-face, there remain several problems with the use of technology to support meetings. Informants in our study identified many issues with the use of teleconferences and video conferences to support meetings over distance and suggested reasons why conducting meetings in virtual worlds would alleviate some of these problems. The potential benefits of using virtual worlds for meetings identified in our interviews include:

Social Engagement: An important research question asks how to bring the social and emotional needs met by face-to-face meetings into meetings over distance (Kahai, Carroll, & Jestice, 2008; Olson & Olson, 2000; Siina, 2007). We learned that our informants feel that virtual worlds provide an opportunity for social engagement, trust, and team building in meetings.

Informal Time Before / After Meetings: One of the ways in which meeting participants build social relationships in meetings is through the ability of meeting participants to socialize and engage in informal discussions before and after the formal meeting time (Olson & Olson, 2000). Our informants indicated that meeting participants in virtual worlds often partake in small group discussions before and after meetings.

In-group / Out-group Effects: Sometimes in meetings over distance, there is one group that meets face-to-face and others who participate remotely. When this happens and the collocated group is larger than the remote groups, there is the potential for the larger collocated group to form a group identify that can be detrimental to the outcomes of the entire meeting group (Bos, Shami, Olson, Cheshin, & Nan, 2004). Our informants indicated that virtual worlds can help overcome the in-group/out-group effects in distributed meetings.

Time Zones and Meeting Times: People we interviewed felt that virtual worlds bring an experience that is closer to video conferencing than teleconferences without forcing people to be on camera in the middle of their night.

Cultural Issues: One of the interesting points that came out in our interviews involves the notion of virtual worlds overcoming cultural issues such that the game-like environment of virtual worlds seems to help people overcome issues of hierarchy prevalent in different cultures.

Getting Attention to Speak or Interrupt: An important barrier to teleconferences is the ability for people to signal their intention to speak or interrupt the current discussion (Isaacs & Tang, 1994). Some of our informants felt that virtual worlds provide an advantage over teleconferences for enabling participants to signal their intention to speak or interrupt the current discussion.

Ability to Re-Experience the Meeting: One of the advantages of technology-supported meetings over meetings in person is the ability to more easily record (and later re-examine) meetings supported by technology (Lucia, Francese, Passero, & Tortora, 2008). Some virtual worlds provide the opportunity to record and later play back the interactions, exchanges, chats, and discussion during sessions and our informants identified this as an advantage over face-to-face meetings.

Multitasking: While the interviewees raised several interesting issues regarding meetings in virtual worlds, the notion of multitasking was the most prominent, mentioned by nearly everyone, often with conflicting opinions. We present our findings about multitasking in virtual worlds meetings next.

5 MULTITASKING IN MEETINGS AND VIRTUAL WORLDS

Multitasking is becoming increasingly common in work environments, partially enabled by technology and ubiquitous computing (González & Mark, 2005; Tang, 2005). In this study, we investigate multitasking in meetings which we characterize as not paying full attention to the meeting agenda. In (Stone, 2006), the concept of continuous partial attention is defined as different from multitasking because multitasking is motivated by productivity and continuous partial attention is motivated by not wanting to miss anything. We do not differentiate among the motivations behind the multitasking. It could take place because one is interested in being more productive (eg., responding to an email while someone makes a point one has heard before) or it could be motivated by not wanting to miss anything (eg., checking for new email or browsing a news source).

Our study finds that this kind of multitasking takes place in meetings most of the time. Almost two-thirds of the people we asked admitted to sometimes, often, or always engaging in multitasking during face-to-face meetings. That increased to 100% in meetings conducted over teleconferences.

It has been argued that multitasking is becoming part of the social norm but, even so, there is a social awkwardness associated with its practice (Tang, 2005). This social awkwardness might be partially explained by awareness and accountability characteristics of socially translucent systems. Erickson and Kellogg (2000) define socially translucent systems as those which support three characteristics of social processes: visibility, awareness, and accountability. Visibility means social information is visible to participants. Awareness means that we are made aware of information that affects our actions hence we bring social rules into play. Accountability means that we know that others are aware that we are there and so we will be held accountable for our actions. If a person is talking with you using a video conference and that person answers their cell phone, you will be aware they are transitioning their attention from you. Since they know you can see them, they will also feel accountable for their actions, possibly causing a level of social awkwardness. If, however, a person is talking with you on a teleconference and their phone is on mute, you may not be aware if they answer their cell phone. More importantly, since they know you cannot see them, they may not feel accountable for the action of answering the cell phone or checking email.

An important question to ask is whether and when this kind of divided attention is useful or detrimental in general work environments and specifically in meeting situations. People have referred to both “Multitasking Attention Deficit” (Curtis, 2000) and, by way of a counter argument, “Multitasking Attention Dexterity” (Torrence, 2001). Tang (2005) suggests that multitasking behavior brings benefit to the individual at the expense of the social good (the experience of the team or outcome of the meeting). He reports that some view multitasking during a meeting as an unwelcome and disruptive activity and others are pleased with the fact that it allows busy executives an opportunity to participate (Tang, 2005).

González and Mark (2005) present a study of 36 information workers in two companies where they found that people continually switch their focus throughout the day, on average spending 10.5 minutes on activities connected to the same work theme before switching to another work theme. Rather than declaring this activity as negative and building systems that make it difficult to multitask, they argue that this kind of multitasking should be supported by collaboration technology and systems so that information workers can do it more effectively.

Specifically to multitasking in meetings, Mark, Grudin, and Poltrock (1999) observed several team meetings conducted using NetMeeting software and found that there were two ways in which people viewed multitasking in meetings. Many people reported that they engaged in other activities during the meeting and most saw this as an advantage of using NetMeeting over face-to-face while others described it as a distraction and a detriment (Mark, Grudin, & Poltrock, 1999).

We asked people at LME to indicate the impact of multitasking on meeting efficiency, effectiveness, outcome, their ability to participate in the meeting, and on the relationships of participants. Results are summarized in Table 1. Just under 74% of the people we asked feel that multitasking has a negative impact on efficiency with almost 16% feeling it has a positive effect on meeting efficiency. Just over 84% feel it has a negative impact on their ability to participate in the meeting. However, only 58% feel it negatively impacts effectiveness with just over 10% stating that multitasking has a positive impact on meeting effectiveness (the rest feel there is no impact on meeting effectiveness). Only 58% of people we asked feel multitasking negatively impacts meeting outcome with just over 5% indicating a positive impact on meeting outcome. Perhaps most surprisingly, just 53% feel that multitasking has a negative impact on participant relationships with the rest indicating no impact on participant relationships.

Table 1. Percentage of people who feel multitasking has a negative impact, no impact, or positive impact on different meeting characteristics.

Meeting Characteristic	Multitasking has a Negative Impact	Multitasking has No Impact	Multitasking has a Positive Impact
Meeting efficiency	73.7%	10.5%	15.8%
Meeting effectiveness	57.9%	31.6%	10.5%
Meeting outcome(s)	57.9%	36.8%	5.3%
Ability to participate	84.2%	15.8%	0%
Participant Relationships	52.6%	47.4%	0%

Our results reinforce the findings of other researchers that multitasking in meetings, while detrimental to some aspects of meetings, is not always negative and can even positively impact meeting effectiveness and efficiency. This seems to indicate that we need technology that allows some level of multitasking to enable its productivity features but not so much as to negatively impact participants’ ability to participate in the meeting. People designing technology to support meetings and collaboration continue to strive to emulate the features of face-to-face meetings with technology (Olson & Olson, 2000). Ideally, we desire a technology that enables a similar amount of multitasking that is afforded in face-to-face environments. In the next section, we argue that virtual worlds provide an environment that is much closer to face-to-face than teleconferencing with respect to multitasking in meetings.

5.1 Multitasking in Virtual Worlds, Face-to-Face, and Teleconferences

Some of our informants described virtual worlds as an ideal environment for meetings when it comes to multitasking but not all agree on the reasons. One felt that it is easier to multitask during meetings in virtual worlds than during videoconference meetings (seeing this as an advantage of virtual worlds), saying, "The other problem is that while you're working in video over IP it's good, but then you're not really multitasking. Whereas when you're on a meeting that's happening in a virtual world, you might be running multiple tasks at the same time. Like running a chat tool in the same environment, listening and watching, talking, moving, you know, through the environment as well as maybe picking up some other things at the same time. This is very hard to do when you're working in a video over IP mode because then people can see that you're not attending the task. And you lose your attention. So those are some of the value points that we're finding these virtual places actually come up over video over IP today."

Another informant felt that multitasking in virtual worlds was difficult, indicating this to be an advantage of virtual worlds for meetings, saying, "I like the productivity [of] meetings because ... it's not a disembodied one off task as you're multitasking ... It's very difficult to multitask when you're ... in virtual worlds ... especially when it's a meeting. So if you don't pay attention your head droops and people know you're not paying attention."

Yet another informant described multitasking during meetings in the following way: "When you're at a meeting you need to be able to sit down and look around the room and pay attention and not stare at your laptop ... on the phone, it's totally opposite and so I'm usually ... wandering around the house and doing the laundry and playing on my computer the whole time. Now the virtual world is sort of *this in between space* and you can be sort of just be doing what you do on the phone, but yet ... you have a lot of multimodal input coming towards you and so you need to not only see what's going on, you need to hear what's going on and you need to follow suit if people are moving around to show that you're sort of paying attention and involved with the conversation or interaction." (Italics added).

We found the idea that virtual worlds were a kind of in between space intriguing and proceeded to examine this notion further by presenting follow-up questions to our informants. We asked them to indicate how often they have multitasked while participating in a meeting in person, using a teleconference, and in a virtual world. We also asked them to indicate how often others were multitasking during a meeting in person, using a teleconference, and in a virtual world. Table 2 summarizes the results.

Table 2. Reporting on how often multitasking takes place in different types of meetings.

Type of Meeting	Rarely / Never Multitask		Sometimes Multitask		Often / Always Multitask	
	Self	Others	Self	Others	Self	Others
Face-to-face	36.8%	10.5%	52.6%	57.9%	10.6%	31.6%
Virtual world	58.3%	36.4%	25%	36.4%	16.7%	27.2%
Teleconference	0%	0%	26.3%	21.1%	73.7%	78.9%

We suggest that a key feature of virtual worlds which may cause this apparent difference is the level of attention required to participate. One informant described it in the following way, "there is a sense that it demands more of your full attention. Because you... have a physical dimension going on. You have to manipulate things to get that there... coordinating, there's often ... the multi channel... dialogues going on, chat, instant message, the phone or the voice… it becomes a slightly more immersive ... experience ... one of the questions is does that bring more accountability."

We further investigated this notion of virtual worlds and engagement by asking people to rate how well face-to-face, virtual worlds, and teleconference meetings enable people to pay attention during the meeting and to what extent they allow people to participate in the meeting. Table 3 shows the percentage of people who said each type supports the two meeting characteristics well or very well. It appears that our informants believe that virtual worlds are better able to support paying attention than teleconferences are but not to the same extent as face-to-face meetings. Our informants also report that virtual worlds provide better opportunity to participate than teleconferences but, again, not as well as face-to-face meetings.

In the next section we discuss our findings and their implications for meetings over distance.

Table 3. Percentage of people who feel different meeting media support the ability to pay attention and opportunity to participate well or very well.

Meeting Characteristic	Face-to-Face	Virtual Worlds	Teleconference
Ability to Pay Attention	94.4%	41.7%	11.1%
Opportunity to participate	77.8%	54.5%	11.1%

6 Discussion

Our study finds that virtual worlds are capable of bringing the distributed meeting experience closer to that of face-to-face meetings when it comes to multitasking. A significant amount of multitasking takes place in meetings, even face-to-face meetings. At the same time, multitasking is seen as both positive and negative. Multitasking can bring benefits such as efficiency and effectiveness but it can also negatively impact the ability of people to participate and pay attention in meetings.

Our study suggests that virtual worlds may be considered as a viable new technological alternative for distributed meetings since they are capable of enhancing the positive aspects of multitasking while at the same time dampening the negative aspects of multitasking. The key features of virtual worlds that enable this balance appear to be the notion of togetherness and presence of individuals through their avatars, which makes people feel they cannot multitask as much as they do in teleconferences, and the level of cognitive attention required to participate in a virtual world (which keeps people engaged in the virtual world meeting) together with the fact that participants are able to switch between the virtual world environment and other technologies. While our informants felt accountable (they do not want their avatar to fall asleep in the meeting), they also felt they can switch between tasks as needed. This notion of accountability may be the key factor that keeps multitasking at bay, preventing it from rising to the levels often seen in teleconferences.

Our study suggests that virtual worlds are a viable tool for distributed meetings. Suggestions for ways to further study this phenomenon are suggested in the following section.

7 CONCLUSION AND FUTURE WORK

The use of virtual worlds in organizations is new and has not yet been studied extensively. In particular, there is little work to date on meetings in virtual worlds. In this paper, we presented results from a study of 23 virtual world users in a large, multinational enterprise. Our study indicates that virtual worlds are a viable IT-enabled environment for meetings over distance and provide an opportunity for multitasking that matches that in face-to-face meetings.

There are several limitations to our study which we discuss here. The number of informants is relatively small and not all of them are experts in virtual worlds nor do all of them use virtual worlds regularly; however, all of our informants worked on at least one virtual world project and some of our informants are experts who not only use virtual worlds for internal business purposes but work with clients as well. The people we interviewed bring a breadth and variability of experience with virtual worlds. Nevertheless, large scale studies or experiments should be conducted to assess the generalization of the findings presented in this study.

Our study indicates that people multitask even in face-to-face meetings. All the informants had participated in meetings where others have multitasked. It is important to understand the impact of this seemingly acceptable behavior on meeting outcomes, relationships among participants, and efficiency. Previous studies and our own results indicate that some people feel the ability to multitask during meetings is beneficial but others find that it negatively impacts several aspects of meetings. This conflicting view of multitasking is an interesting avenue for future research.

It would be interesting to explore multitasking in different technology-enabled meetings and examine its effects on various outcomes (e.g., problem solving and decision making). Another promising direction is to study participants' satisfaction with the meeting outcomes and relate it to their own level of multitasking, and that of other participants. These studies could help us learn if there are unique ways in which meetings should be facilitated when it comes to different technologies and multitasking.

REFERENCES

[1] Balkin, J.M. & Noveck, B.S. (2006). Eds., The State of Play: Law, Games, and Virtual Worlds, New York University Press, New York, 2006.

[2] Barnes, S. (2007).Virtual Worlds as a Medium for Advertising. Database for Advances in Information Systems 38(4) 45-55.

[3] Bessière, K., Ellis, J. B., & Kellogg, W. A. (2009). "Acquiring a Professional 'Second Life:' Problems and Prospects for the Use of Virtual Worlds in Business", CHI 2009, April 4-9, 2009, Boston, Massachusetts, USA, ACM 978-1-60558-247-4/09/04, 2883—2898.

[4] Bos, N., Shami, N. S., Olson, J. S., Cheshin, A., & Nan, N. (2004). In-group/Out-group Effects in Distributed Teams: An Experimental Simulation, Computer Supported Collaborative Work, 2004, 429—436.

[5] Boulos, K.M.N., Lee, H., & Wheeler, S. (2007). Second Life: An Overview of the Potential of 3-D Virtual Worlds in Medical and Health Education. Health Information and Libraries Journal 24 233–245.

[6] Bronack, S., Riedl, R., Tashner, J., & Greene, M. (2006). Learning in the Zone: A social constructivist framework for distance education in a 3D virtual world. Proceedings of Society for Information Technology and Teacher Education International Conference, 2006, 268-275.

[7] Cameron, A. F. & Webster, J. (2005). "Unintended consequences of emerging communication technologies: Instant messaging in the workplace," Comput. Human Behavior, vol. 21, no. 1, pp. 85–103, 2005.

[8] Curtis, H. (2000). Flash Web Design: the art of motion graphics, New Riders Publishing, 2000.

[9] Edwards, C. (2006). Another world. Engineering and Technology Dec. 2006 28-32.

[10] Erickson, T. & Kellogg, W. A. (2000). Social translucence: An approach to designing systems that support social processes. ACM Transactions on Computer-Human Interactions, 7, 1 (Mar. 2000), 59-83.

[11] Goel L. & Mousavidin, E. (2007). vCRM: Virtual Customer Relationship Management. Database for Advances in Information Systems 38(4) 56-60.

[12] González, V. M. & Mark, G. (2005). Managing currents of work: Multi-tasking among multiple collaborations, In Proceedings of the 8th European Conference of Computer-supported Cooperative Work (ECSCW'05), September 18-22, 2005, Paris, France.

[13] Hemp, P. (2006). Avatar-Based Marketing, Harvard Business Review 84(6) 48-56.

[14] Hendaoui, A., Limayem, M., & Thompson, C. W. (2008). 3D Social Virtual Worlds: Research Issues and Challenges. IEEE Internet Computing 12, 1 (2008), 88-92.

[15] Holzwarth, M., Janiszewski, C., & Neumann, M. M. (2006). The Influence of Avatars on Online Consumer Shopping Behavior, Journal of Marketing 70 (4) pages 19-36, 2006.

[16] Isaacs, E. A., & Tang, J. C. (1994). What video can and cannot do for collaboration: A case study. Multimedia Systems, 2, 63–73.

[17] Kadavasal, M.D., Dhara, K.K., Wu, X. and Krishnaswamy, V. (2007). Mixed reality for enhancing business communications using virtual worlds, Proceedings of the 2007 ACM symposium on Virtual reality software and technology 233-234.

[18] Johnson, B.C., Manyika, J.M., & Yee L.A. (2005). "The next revolution in interactions", The McKinsey Quarterly, 2005.

[19] Kahai, S. S., Carroll, E., & Jestice, R. (2007). Team collaboration in virtual worlds. ACM SIGMIS Database, 38(4), 61-68.

[20] Linden Lab. (2009). Case Study: How Meeting in Second Life Transformed IBM's Technology Elite into Virtual World Believers, Accessed August 29, 2009, available at http://secondlifegrid.net.s3.amazonaws.com/docs/Second_Life_Case_IBM_EN.pdf

[21] Lucia, A. D., Francese, R., Passero, I., &Tortora, G. (2008). SLMeeting: supporting collaborative work in Second Life. In Proc. AVI 2008, ACM Press (2008), 301-304

[22] Mackinlay, J. D., Card, S. K., & Robertson, G. G. (1990). Rapid controlled movement through a virtual 3D workspace. SIGGRAPH Computer Graphics, 24, 4, 1990, 171-176.

[23] Mark, G., Grudin, J., & Poltrock, S. E. (1999). "Meeting at the Desktop: An Empirical Study of Virtually Collocated Teams", In Proceedings of ECSCW'99, The 6th European Conference on Computer Supported Cooperative Work, 12-16 September 1999, Copenhagen, Denmark, pp. 159-178.

[24] Mayer-Schoenberger, V. & Crowley, J. (2006). Napster's Second Life? The regulatory Challengers of Virtual Worlds, Northwestern University Law Review 100(4) (2006) 1775—1826.

[25] Messinger, P. R., Stroulia, E., Lyons, K., Bone, M., Niu, A., Smirnov, K., & Perelgut, S. (2009). "New Directions for Social Computing in Virtual Worlds: Applications for

Business and Social Sciences," Decision Support Systems, Vol. 47, Issue 3, 204—228.

[26] Nebolsky, C., Yee, N. K., Petrushin, V. A., & Gershman, A. V. (2003). Using virtual worlds for corporate training, Proceedings of the 3rd IEEE International Conference on Advanced Learning Technologies (ICALT'03)

[27] Olson, G. M. & Olson, J. S. (2000). Distance Matters, Human-Computer Interaction, 2000, Volume 15, pp. 139–178

[28] Siina, R. M. (2007). Emotional Engagement on Geographically Distributed Teams: Exploring Interaction Challenges in Mediated Versus Face-to-Face Meetings, PhD Dissertation, Stanford University, June 2007, 174 pages.

[29] Sipress, A. (2006). Where Real Money Meets Virtual Reality, The Jury Is Still Out, Washington Post (Dec. 2006) Retrieved August 12, 2009, from http://www.washingtonpost.com/wp-dyn/content/article/2006/12/25/AR2006122500635.html

[30] Soraker, J. H. (2008). Global Freedom of Expression Within Nontextual Frameworks, Information Society 24(1) (2008) 40—46.

[31] Stone, L. (2008). Continuous Partial Attention -- Not the Same as Multitasking, Business Week, July 2008, Accessed July 3, 2009, available at http://www.businessweek.com/business_at_work/time_management/archives/2008/07/continuous_part.html

[32] Strauss, A. & Corbin, J. (1994). Grounded Theory Methodology: An Overview. In N. K. Denzin, & Y. S. Lincoln (Eds), Handbook of Qualitative Research. Thousand Oaks, CA, US: Sage Publications, Inc. 273-285.

[33] Tang, J. C. 2005. "Ubiquitous computing: Individual productivity at the expense of social good?", First International Workshop on Social Implications of Ubiquitous Computing, held at CHI 2005, April 2005, Accessed July 3, 2009, available at http://www.vs.inf.ethz.ch/events/ubisoc2005/UbiSoc%202005%20submissions/04-Tang-John-NEW.pdf

[34] Torrence, P. (2001). "Fast Cars, Fast Food, Fast Access", Digital Web Magazine, September 9, 2001, Accessed September 1, 2009, available at http://www.digital-web.com/articles/fast_cars_fast_food_fast_access/

[35] Yee, N., Bailenson, J., Urbanek, M., Chang F., & Merget, D. (2007). The Unbearable Likeness of Being Digital: The Persistence of Nonverbal Social Norms in Online Virtual Environments, CyberPsychology & Behavior 10(1) (2007) 115—121.

Towards Smart Privacy on the Personal Web

Reza Samavi, Mariano P. Consens

{samavi, consens}@mie.utoronto.ca, MIE, University of Toronto

Abstract

User-centric privacy management is an important component of the Personal Web, and even more so in the context of personal health applications.

In this position paper, we propose a logical framework for *smart privacy* based on a modular and extensible ontology that supports reasoning about privacy from a very broad range of perspectives. We describe the motivations behind the development of smart privacy and outline key features of the logical framework in the context of Personal Health Record applications.

1 Introduction

Personal Web is an emerging research topic driven by the transformation of the Internet and web from the way users currently interact and navigate resources on the web, to a smart paradigm mainly centered on users' experience. The main goal of smart Internet is to empower user, as an individual, to seamlessly and smartly self-manage the vast amount of web resources and services to achieve her personal goals [4]. A user-centric perspective of service utilization requires users to play an active role in the process. The promise of personal web is to make this shift socially and cognitively viable. Such a perspective brings new challenges in design and architecture of the web. For example, how to enable the underlying web architecture to support and deliver services based on the individual needs; users need to be able to customize the process according to the task they as individuals want to accomplish, and in such a way that respects their preferences. One such an imperative task is the self-management of privacy.

To meet the promises made by *Smart Internet* [4] and to make the personal web a real user-sovereign system, we propose a notion of *smart privacy* in which the essence of privacy is embodied in system design. We argue that a semantic flow model [2] can provide an unambiguous, and computer-interpretable descriptions of information flow and its relevant norms on personal web. We adopt the point of view of privacy as contextual integrity [13], and propose an end-to-end privacy model in which rather than focusing on sensitivity of data or roles, a permutation of semantically rich contextual information provides users with support to self-manage their privacy.

In this position paper, we make the case for Personal Health Records (PHRs) as an emerging personal web application that can play a key role in transforming different health care processes. From the privacy perspective of personal web, this is a perfect candidate, since the most sensitive information about one is at stake.

In the following subsections we discuss the motivations and expected contributions of this research alongside a brief overview on the state of privacy frameworks on web. In Section 2, we motivate the needs for the model with a PHR use case. Section 3 outlines the proposed privacy model based on the novel notion of semantic flow model. Finally, section 4 concludes the paper with a few remarks on future directions.

Motivation. PHRs are the type of electronic records managed by patients that may end up becoming the least fragmented picture of an individual's health [17]. PHRs, from their first appearance nearly a decade ago as being a patient-centric repository of health data, have been gradually transformed into open software platforms with published application programming interfaces (APIs) [10]. This is similar to what we are experiencing with other emerging personal information platforms like *Facebook*. With these APIs, PHRs are now viewed as *personal health information system* that enables aggregation of different types of health data across multiple healthcare providers [10]. Furthermore, it allows a growing number of innovative third-party applications to emerge around PHRs. Applications for online-health assessment, for participating in clinical trials, or for finding low cost medicine are just a few of the examples. Some futuristic application scenarios of PHRs are also emerging from Web 2.0 communities when the power of social networking combines with patients' empowerment enabled by PHRs [16] . Studies show that patients' trust on the information received from peers with similar conditions is surpassing the trust in doctors and academic experts. This indicates the powerful role of social networking in public health with the intervention of PHRs [6].

While PHR platforms have great potential for patient empowerment, if the consequences of sharing and data usage are not clear to patients they are prone to greater risk of privacy breaches. With patients having complete control over their health information, the central issue to PHR becomes the self-administration of privacy by patients; how patients, *themselves,* can easily and effectively control their privacy, while sharing their health profile.

Existing privacy architectures for PHRs are rather primitive and the consent to share the record typically applies to the entire record. The patients' administrative role in controlling their privacy is limited to a trivial binary consent of "I agree" or "I disagree". More importantly, there is not mechanism in place for controlling the *usage* of data and understanding the consequences of sharing. Controlling usage after second and further sharing, i.e., when a receiver shares a patient's record with another party, is even more challenging.

The research outlined in this paper aims to fill these practical gaps and propose a logical framework for expressing patient's privacy expectations and mechanisms to ensure these expectations are enforceable. We believe an authenticate design of privacy system for a personal web application will be successful only if the system is semantically sound. The system should precisely capture what the user expectations are and how they affect occurring of *future* activities in a process in order to access to health information occurring *now* (obligation and promise management). Therefore, we use first order logic formulas for expressing and reasoning about these actions and norms that govern flow of information from one application to another in the PHR. We construct our logical model based on the framework of *contextual integrity* [13]. In contrast to the classical view of privacy that mainly focuses on type of data or role of objects, contextual integrity provides a systematic guidance to incorporate all elements of context (the actor, role, purpose, type of data, etc.) alongside with the norms of transmission for information.

In this research, we propose the logical framework for smart privacy as a modular and extensible ontology. This ontology is intended to support reasoning about privacy from a very broad range of perspectives, e.g. privacy preferences, third party privacy policies, delegation of privacy policies, and obligation enforcement. Instead of designing the logic from the scratch, we build the ontology using ISO 18629, Process Specification Language (PSL)[2].

Related Work. A number of frameworks for defining and enforcing privacy policies have been proposed, including P3P [5], XACML [14], EPAL [9], LPU [1], Privacy APIs [11]. These existing solutions cannot adequately address personal web privacy management requirements for the following reasons. First, these frameworks are mainly designed having institutional privacy needs in mind. In other words they are built to protect organizations from being liable in case of breaching privacy laws and regulations, not for the purpose of supporting an individual who has personal privacy preferences. Second, these frameworks either are not expressive enough to support some important features of privacy obligations such as repeating obligations, and multiple responsible agents [14], or use complex logical machinery that makes practical usage of the framework infeasible (e.g. [1]). In contrast, the approach suggested in this research by exploiting first order theories in PSL is highly expressive, while PSL constructs also can be easily and systematically extended.

Contributions. We make three contributions. First, we contribute to a specific personal web domain, i.e. PHR, by enabling PHR consumers to effectively express their privacy expectations and define their social privacy boundaries associated with their health information.

Second we design a novel logical privacy framework in which *user preferences* (in contrast to laws and regulations) form the foundation of the privacy management. These preferences dynamically regulate the flow of information in a privacy-sensitive process. The privacy constraints in the model are introduced using the same semantic constructs used to express all other process constraints (for example task ordering, concurrency, task decomposition). This allows achieving privacy goals and utility goals of a process in positive-sum paradigm as described in Privacy by Design philosophy [3].

Third, our work provides a novel ontology-based obligation model based on PSL to the privacy modeling communities in general. This approach can be applied in other domains of personal web such as social networking.

2 Use Case and Desiderata

We illustrate with an example the need for semantically rich and computer-interpretable behaviour abstraction for process model in a personal web domain. In a simple communication between a hypothetical PHR user, Alice and a third-party service provider, MedicaSave, Alice wants to link her prescription information with the service. MedicaSave reads patient's prescription data stored in PHR and finds the generic equivalent of the drug and low-cost pharmacies for the prescribed drugs. When Alice selected which pharmacy she prefers, then MedicaSave shares Alice's personal information with the pharmacy to deliver the medicine. The Pharmacy needs to access Alice's PHR to check a potential drug interaction. An excerpt from Alice's privacy preferences is shown below:

1. I do not want my sensitive health information to be released to any other third parties except those who are in my circle of care.
2. My explicit consent is required only when the entity using my data is not a covered entity.
3. I would like to be notified monthly about the status of my data and every time my information has been viewed not later than 1 day after access.
4. I would like to receive my health relevant information by email.
5. My information on third party services should be deleted when I unlink my PHR profile to the service.

In the other side of the disclosure, MedicaSave and Pharmacy have also their own privacy policies in a form of a document that describes how they handle private information. The privacy system should be able to address the following set of high-level queries:

Decision making queries: Should Alice agree with the MedicaSave privacy agreement? In what condition is MedicaSave allowed to send Alice's information to Pharmacy? Does pharmacy need Alice's explicit consent to access her PHR?

Contextual queries: Is MedicaSave a covered entity? Is the Pharmacy a member of Alice's circle of care? Is the prescription information a sensitive health data of user? Does sending an email by an agent q counts as a notification?

Obligations and Provisions queries: In a given time point t, what type of notifications need to be sent by MedicaSave and/or Pharmacy? In a given time point t, should Alice's data be retained?

Exception queries: Can data item d be still accessed by agent p if the notification has not been sent?

As this simple example shows, answering these queries requires a computer-interpretable process model with rigorous and complete axiomatization that support automated reasoning about the concepts. In [15] authors showed how privacy model requirements can be extracted from scenarios, our investigation of more than twenty PHR usage scenarios (range from standard clinician- patient scenarios to futuristic usage of PHR data across social networks) led us to the following high-level desiderata for a smart privacy model.

Privacy preferences: We distinguish between two sides of disclosure, user and service. As noted above with user-centric approach, smart privacy should be able to capture first the flexibility which is inherent in *preferences* and second the privacy notions that might be of interest to a user but not specified in laws and regulations, for example concepts such as appropriateness, embarrassment, reciprocity, and deservedness.

Generic privacy concerns: Each individual service offers different privacy settings and agreements. Therefore, while privacy concerns of a user has not been changed, user is required to endlessly repeat many decisions to perform a single task of expressing privacy expectations. The privacy framework must be able to communicate the semantics of an individual's privacy preferences with multiple services.

Semantic support for privacy concepts: Users of the application are diverse. They come from different jurisdictions and with different cultural backgrounds. Therefore, the privacy framework should support different terminologies of concepts.

Granularity of privacy concerns: Privacy is a subjective property. Users of an application span from those who believe "privacy kills, openness heals" to those with very strict privacy concerns. The privacy framework should be able to capture this diversity.

Support privacy as a process: Different from access control mechanisms that an action is allowed or denied for a specific role in a specific point of time, privacy control mechanisms in general, and in Personal Web applications in particular, require reasoning about a process, a series of temporally constrained actions and occurrences.

Support temporal constraints: Reasoning about temporal constraints is essential in a privacy framework [8]. There are actions need to be performed on data objects before an access is authorized (Conditions), and/or, actions need to be performed after an access is authorized (obligations).

Support multiple responsible agents: In a distributed computing system, not always the agent responsible for an obligation activity is necessarily the same agent who performed the action caused the obligation. Therefore the logical framework must be able to reason about the agents responsible for actions.

Support context representation: In privacy as contextual integrity paradigm, transmission norms, largely depend on the context in which the communication activities occur. The privacy framework should reason about the characteristics of the context.

Support expressing of utility goals: The logical framework must be able to capture the logic of utility and privacy together [1]. In many cases reasoning about the privacy actions requires reasoning about the goal of the process itself. A privacy framework cannot reason about the privacy goals while precluding reasoning about the activities required the goal of the workflow to be achieved.

Usability for self-administration of privacy: users are neither policy makers nor system administrators; therefore, the system must be such usable that an average user can express her privacy expectations. In other words system support is required to improve comprehensibility and consciousness of an average user on controlling her privacy.

Although some of these requirements include elements of a user interface design, our focus in this research is on the infrastructure that is needed to support the required functionality. Nevertheless, the characteristics of the user interface substantially influence how the user can interact with a system.

3 Smart Privacy Model

In this research, we propose the logical framework for smart privacy as a modular and extensible ontology. This ontology is intended to support reasoning about privacy from a very broad range of perspectives, such as user privacy preferences, third party privacy policies, delegation of privacy policies, and obligation enforcement. Instead of designing the logic from the scratch, we build the ontology using ISO 18629 (Process Specification Language (PSL)[2]) upper ontology. Upper ontology is an ontology, which describes very general concepts that are the same across all domains.

In smart privacy, theories of PSL are extended to express specific privacy constructs such as pre-conditions, post-conditions, obligations, and communication behaviors as constraints over occurrences of process activities in the context of a rigorously axiomatized first-order logic framework (Figure 1). The brief review of contextual integrity reveals how the ontologies are used to gain insight with regard to privacy.

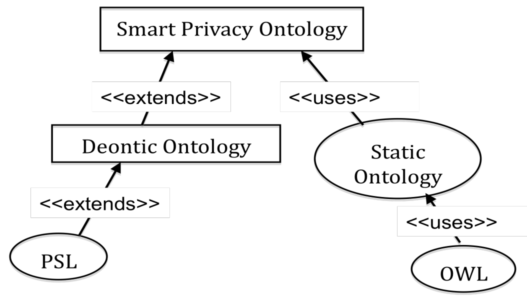

Figure 1: Ontology-based Smart privacy

Contextual integrity. Contextual integrity is a philosophical account of privacy that provides a normative model, or framework, for evaluating the flow of information between agents [13]. In contrast to the classical view of privacy as "control over information about oneself"[18], contextual integrity emphasizes on evaluating why certain patterns of flow of information provoke the sense of privacy violation, while others not. The model then defines the notion of *appropriateness* to answer this question. Disclosing information per se is not what makes us feel breach of privacy. We feel our privacy is violated when the same information communicates in a context that we feel it is embarrassing or inappropriate. For example sharing information about one's unwanted pregnancy to her family physician looks fine and appropriate while sharing the same information to her boss causes one's outcry. Therefore, features of contexts accompanied with the flow of information determine privacy. Five constructs are keys to defining contextual integrity: contexts, informational norms, actors, attributes, and transmission principles. Contexts are structured social settings characterized by the roles that actors play, by certain ends or values that a context forms around it (e.g. the value of a health care context is providing health service), and the norms that prescribe and proscribe acceptable actions and practices [13, p133]. Attributes define the nature of the information in question. If agents refrain to share any attribute, there would not be a violation of privacy. Transmission principle is a set of constraints that governs the information flow. Norms prescribe which transmission principles ought to govern the flow of information and privacy violation arises if these principles are not followed.

In our model we formalize the main concepts of contextual integrity using ontologies as shown in Figure 1. Our model consists of communicating agents (G) who take various roles

(*R*). Each role (*r*) has certain capacities (*U*). These capacities capture knowledge about the ability of an agent to pursue one or more utility goals. Agents participate in Communication Activities (*A*). For each communication activity, at least three distinct type of agents are involved; the principal agent whose information is in stake; the sender of the information; and the receiver of the information. Each receiver agent has a set of history of access to information as its property. Occurrence of each activity changes the state of an agent in terms of the history of access. Bind to each history, there exists activities that must occur and/or activities that must not occur. Occurrences of these activities are either temporally or causally constrained and are accompanied by norms of transmission specific to each context. This specification can be modeled using ontologies.

Deontic Ontology. Our semantics are based on the concept of provisions and obligation linked to the PSL *activities*. Predicates that change due to activity-occurrences are modeled using *fluents*. Pre-conditions and post-conditions modeled as activity-occurrences with temporal properties of start-point and end-point. In this way we can entail whether a privacy constraint (privacy goal) is *successful*, *failed*, or *violated;* utility goal is achieved or failed. The construct of occurrence tree in PSL captures what we call privacy contingency plans for either of these conditions. Set of all possible activity-occurrences of a process is modeled as a tree whose nodes correspond to individual activity-occurrences and where the children of one activity-occurrence correspond to the set of all possible activity-occurrences that could immediately follow it. For a fluent such as *send_notification(t, p, q)* in time *t* from agent *p* to agent *q* and *access* occurrence of *o* the value of the fluent just prior to *o* occurs is expressed by *prior(send_notification(t, p, q), o)* and the value of the fluent just after the *o* occurs is expressed by *holds(send_notification(t, p, q), o)*. By using these basic constructs, more complex behaviour execution constructs such as if-then-else can be modeled in PSL.

Deontic ontology in our logical framework captures the generic knowledge about a context and its norms of transmission. Both classes (contexts and norms) are defined in terms of activities and their occurrences using first order formulas. The novelty of our work is that, in contrast to the previous formal privacy frameworks (e.g. [1, 12]), *context* in our model is not expressed as an entity by itself, but it comes to the picture by a specific permutation of agents, roles, purposes, environmental variables and activities. Then a set of deontic constraints defines the properties and relationships of these entities along with *norms* of the context. This approach captures dynamic and extensible nature of context in our privacy model.

The second class of axioms in deontic ontology represents the *transmission norms* that govern a context's information flow. For example, the constraints such as "I would like to be notified by email every time my information has been accessed" will be presented as a deontic constraint:

for all (o1) (implies (occurance_of (o1, record access)), (exists (o2) (and (occurance_of (o2, send_email), (begine_of (o2) > (begine_of(o1))))))

We use the same semantic as the one we used for contexts to represent these principles. By this approach, while evaluating a privacy policy, the reasoning engine collects requirements from *all* applicable axioms, which restrict a *context* and its *norms,* in one step. This allows addressing an important property of contextual integrity. We can now, relate the issues of compliance and refinement of privacy policies to the logical concepts of satisfiability and entailment. Entailment is key to understand whether a context and its norms, comply with the transmission principles. Our semantic also explicitly captures the concept of past and future. Thus, reasoning is not limited to the time point of the access, but the decision of compliance can depend on what actions have occurred previously and can require occurrence of future actions. The deontic ontology in the proposed logical framework consists of two sets of first order axioms: Contexts ontology, and transmission norms ontology.

We view agent communication in PHR context as processes with specific goals. Therefore, there are many common concepts of process model in general that we are going to solicit from PSL for reasoning about a context and its privacy norms. Concepts such as *before, after, activity, activity occurrence,* are examples of some common terms that are already expressed by PSL ontology.

Static Ontology. The deontic ontology as described above works in the spirit of a static ontology. The static ontology in our model characterizes classes of entities used in deontic ontology, their properties and their relationships to each other. For example, unambiguous definition of entities such as *Data_receiver, Covered_entity, Uncovered_entity* is essential for reasoning about privacy constraints. Agents in our logical model of privacy are autonomous, so the same constraint across multiple agents may be stated differently. The static ontology contributes to our logical model by providing support for interoperability and more effective use of knowledge about contexts and their information transmission norms. Furthermore, it supports neutral authoring of privacy constraints [7] in a sense that user privacy preferences can be authored in a single language with the service privacy policies. This static ontology will be formulated in description logic (DL), supported by the web ontology language OWL DL [12]. Figure 2 shows a partial representation of norms static ontology.

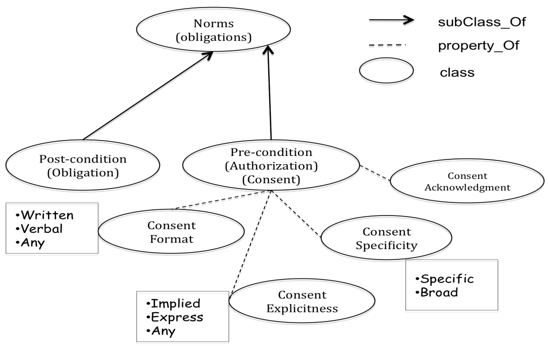

Figure 2. Static ontology (partial) for Norms

4 Closing Remarks

In this paper we introduced the concept of smart privacy for personal web. We described the motivations behind smart privacy, alongside the features that a logical model needs to adhere in order to adequately express privacy requirements of an important class of personal web applications: PHR.

We argued that existing privacy languages and frameworks are not expressive enough to capture all aspects of smart privacy. As such, we rationalized using formal ontology language to create an extensible semantic flow model for user-centric privacy management in personal web. Although this logical model per se cannot address all requirements of smart privacy mentioned in this paper, it provides a solid semantic foundation. How this logical system can architecturally be materialized, and how a user can communicate with this logical system, are unanswered questions and part of our future research.

Acknowledgements

Special thanks go to Dr. Thodoros Topaloglou for his helpful comments on the problem formulation, and Prof. Michael Gruninger for his valuable help in the designing of ontology-based privacy model. Financial support from the Natural Sciences and Engineering Research Council of Canada and IBM Privacy Award are greatly acknowledged.

References

[1] A. Barth. Design and Analysis of Privacy Policies. *Phd thesis*, Stanford U. 2008.

[2] C. Bock, M. Gruninger, PSL: A semantic domain for flow models, *Soft. and Sys.Modeling* 4:209-231. 2005.

[3] A. Cavoukian. Privacy By Design, Take the challenge. *Office of information and privacy Commissioner of Ontario*, Canada, 2009.

[4] J.R. Cordy, M. Chignell, J. Ng. SITCON:The CAS/NSERC Strategic Workshop in Smart Internet *Tech. Proc. of the CASCON 2009. Center for* Advanced Studie*s o*n Collaborative Research. Nov. 2009.

[5] L. F. Cranor, M. Langheinrich, M. Marchiori, M. Presler- Marshall, and J. R*eagle. The platform for pri-*vacy preferences 1.0 (P3P1.0) specification, 2002. http://www.w3.org/TR/P3P/

[6] Edleman Trust Barometer. http://www.edelman.com/trust/2008/. 2008

[7] M. Gruninger, The Ontological Stance for a Manufacturing Scenario, to appear in Journal of Cases in Information Systems. 2009.

[8] M. Hilty, D. A. Basin, A. Pretschner, On Obligations, *in: Proc. of ESORICS'05*, Springer-.pp. 98{117. 2005.

[9] G. Karjoth, M. Schunter, M. Waidner, Platform for Enterprise Privacy Practices: Privacy-enabled Management of Customer Data, in: *Proc. Of PET'02*, Vol. 2482 of LNCS, Springer-Verlag. pp. 69{84. 2002

[10] K. D. Mandl. W. Simons. W. Crawford, and J. Abbett, Indivo: a personally controlled health record for health information exchange and communication. *BMC Med Inform Decis Mak*. vol. 25. 2007.

[11] M.J. May, C.A. Gunter, and I. Lee. Privacy apis: Access control techniques to analyze and verify legal privacy policies. *In IEEE Workshop on Computer Security Found*, pp 85{97. IEEE Comp Society, 2006.

[12] B. Motik. P. Patel-Schneider. B.C. Grau. OWL 2 Web Ont. Lang.:Direct Semantics. Tech. rep W3C 2008

[13] H. Nissenbaum. Privacy in Context: Technology, Policy, and the Integrity of Social Life. Chicago: Stanford Law Books, 2009.

[14] OASIS, eXtensible Access Control Markup Language (XACML) Version 2.0, OASIS Standard (2005). URL http://docs.oasis-open.org/xacml/2.0/access control-xacml-2.0-core-spec-os.pdf

[15] R. Samavi, and T. Topaloglou, Designing Privacy-Aware Personal Health Record Systems, *Proc. of the ER 2008 (CMLSA)*, Barcelona, Spain, 2008.

[16] Stewart, Darin. " EcontentMag.com: Socialized Medicine: How Personal Health Records and Social Networks Are Changing Healthcare." http://www.econtentmag.com/Articles/ArticleReader.aspx?ArticleID=56166&PageNum=2, (acc. 13, 2010).

[17] P.C. Tang, D. Lansky. The missing link: bridging the patient-provider health information gap. *Health Aff*. vol. 24. pp 1290-1295. 2005

[18] A.F. Westin. Privacy and freedom, New York, Atheneum, 1967

Personalized Services

A Framework for Composing Web Resources

Hua Xiao
School of Computing
Queen's University
Kingston, Ontario, Canada
huaxiao@cs.queensu.ca

Bipin Upadhyaya, Ran Tang, Ying Zou
Dept. of Electrical and Computer Engineering
Queen's University
Kingston, Ontario, Canada
{9bu, ran.tang, ying.zou@queensu.ca}

Joanna Ng, Alex Lau
IBM Toronto Lab
Markham, Ontario, Canada
{jwng, alexlau}
@ca.ibm.com

Abstract

Large amount of heterogeneous Web resources, such as SOAP-based Web Service and RESTful service exist on the Internet. It is labor-intensive and inefficient for an end-user to search and compose different Web resources in order to fulfill his/her requirement. To support the end-user's various activities, we propose a Web resources composition framework. This framework can help end-user: 1) discover available Web resources to fulfill the end-user's goal, despite of their types; 2) represent the relation between different resources to allow them to be used collaboratively; 3) automatically compose required Web resources to fulfill the goal specified by the end-user.

1 Introduction

Various types of Web resources, such as SOAP-based Web Service and RESTful service, exist on the Web to provide various functionalities, such as information access and online banking. There are large amount of Web resources available on the Internet. Not all Web resources are highly relevant to a user's requirement. Only a subset of Web resources can fulfill an end-user's requirement. It is difficult for the Web end-users to sift through the sheer volume of Web resources to fulfill their goals. Without the aid from a service composition tooling, an end-user has to manually discover and compose different Web resources. For example, a person planning a conference trip needs to locate the Web resources for transportation, accommodation and other activities separately and integrate the results from these Web resources. This is a time-consuming and tedious process and may not produce the optimal outcome. For example, the end-user may not be able to discover the Web resource that provides the most economical air ticket.

To support the end-user's social, professional, recreational and other activities, it is essential to create a Web-based service composition framework that can 1) discover all available Web resources to fulfill the end-user's goal, despite of their types; 2) represent the relation between different resources to allow them to be used collaboratively; 3) automatically compose required Web resources to fulfill the goal specified by the end-user. It is challenging to realize such a service composition framework. In particular, The Web resources are described in heterogeneous formats. For example, Web Service Description Language (WSDL) [2] is used to describe SOAP (Simple Object Access Protocol) based Web Services that makes remote procedure calls. HTTP-based APIs are simpler Web resources, implemented as a set of standard HTTP requests. Examples of HTTP-based APIs are twitter, Flickr and various Yahoo APIs. The HTTP-based APIs can be described using Web pages or WSDL 2.0. Informational Websites are implemented by various technologies, such as Ajax, HTML and XML. Some of the descriptions, such as HTTP-based APIs are not machine readable. This hinders the ability to discover the most suitable resource to satisfy the goal of a particular end-user.

In this paper, we propose a framework to help end-users compose various Web resources. Our framework uses a unified description schema to describe the heterogeneous Web resources and a resource graph model to represent the relations among different Web resources. Representati tional State Transfer (REST) [9] is an architectural style for network-based systems. REST was not introduced as an approach to designing Web services, yet the non-corporate Web Service community as alternative to SOAP/WSDL has adopted it. Although not always adhering to the all of REST's constraints [10], RESTful Web Services are gaining popularity and are adopted

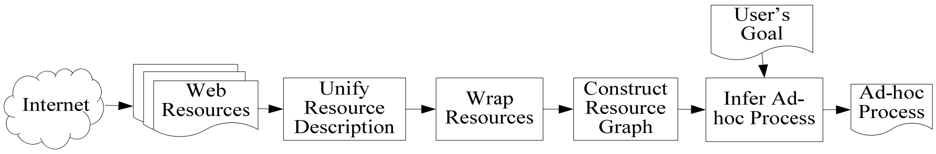

Figure 1: Steps for composing Web resource

by major service providers like Google, Amazon and Yahoo. RESTful service provides uniform interface which is immutable (no problem of breaking clients). HTTP/POX is ubiquitous (goes through firewalls). Since it adheres to the principle of Web, it naturally has proven scalability with caching, clustered server farms for Quality of Service (QoS). End user just need browser to get started, no need to buy WS-* middleware. Moreover, we provide an approach to compose resources for end-users using the resource graph. The composite resources are represented as ad-hoc processes. An ad-hoc process contains a set of tasks without strict execution order.

The remainder of this paper is presented as follows. Section 2 gives an overview of the proposed framework. The details of the unified description schema, the resource graph to model the relation among resources, and the technique to construct ad-hoc process using the resource graph are presented in section 2. Section 3 concludes this paper.

2 Overview of Framework

Figure 1 provides an overview of our framework. To describe the heterogeneous Web resources, we collect heterogeneous Web resources from the Internet and represent them in a unified resource description schema. We identify the relations among different resources and construct a resource graph. HTTP-based APIs are generally described using plain, unstructured HTML documents which are only useful to human developers. Nowadays, using Web resources, such as finding suitable services, composing services, mediating between different data formats, are mainly manual tasks. To maximize the interoperability among the resources, we need a common data model to describe the resources and their relations. There are two requirements for interoperability: (1) the Web resources themselves must be able to programmatically interoperate. For example, they must be able to invoke one another and pass data among themselves; (2) there must be a user interface mechanism for the user to orchestrate the Web resources to work together toward some complex goal. End-users should be able to compose and define the flow between the Web resources. RDF is designed specifically for exchanging and integrating Web data. In our framework, we wrap the unified Web resources into RESTful services then adopt Description Framework (RDF) [1] to describe RESTful services and their relations. While we construct the resource graph, the unified resource descriptions are used to help us identify resources and their relations.

When an end-user wants to fulfill a goal, the end-user simply describes the desired goal using keywords. We map the keywords into resources described by the resource graph. and infer an ad-hoc process from the resource graph to help the end-user fulfill the goal. In the following subsections, we discuss the details in unifying the description of various types of resources, constructing a resource graph, and inferring ad-hoc processes from the resource graph.

2.1 Unifying Resource Description

To assist the automatic discovery of various Web resources, we propose a schema to uniformly describe different types of Web resources. The unified representation provides a better chance to discover Web resources than limiting the service discovery within the Web resources of a single type.

As shown in Figure 2, we define two parts in the unified schema: the general description part and the operation description part.

The general description of a Web resource provides a bibliographic description about the Web resource: the type, the name, the provider and the URI of the Web resource. Such descriptions are common to all types of Web resources. There are standards suitable for representing the general description. For example, using 15 text fields (e.g., title, type and publisher), the Dublin Core metadata schema can describe various resources, e.g., books and Web pages [3]. We adopt the Dublin Core format to represent the general description.

Among the four fields of general description, the type, the name, the provider fields are self-descriptive. The URI field of the general description refers to the URI that identifies the Web resource on the Internet. The URI of a SOAP-based Web Service points to its WSDL file. For an HTTP-based API, the URI field is filled by the URI of its description Web page. To invoke a particular functionality of the Web resource, another URI may be required because each operation of the Web resource may have a different URI, which is described in operation description part.

The operation description describes the functionalities offered by a Web resource. A Web resource can deliver one or more functionalities. An operation represents a primitive unit of functionality used to compose a service-oriented application. Different types of Web resources contain varied number of operations. Most SOAP-based Web Services and HTTP-based APIs provide complex functionalities and contain multiple operations.

To describe each operation, we use the tag-based description, the formal interface and the excerpt of existing description. The tag-based description uses a set of descriptive tags (i.e., keywords) to informally represent an operation, including the functionality description, the input description and the output description. The input/output descriptions help describe different operations with the same name or similar functionality description. For example, two operations are named as "displayOrders". One operation with the parameter "productID" is different from the other one with the parameter "custormerID". The tag-based description can concisely convey the functionality of the operation to the resource consumers. In addition, it facilitates automatically compare the functionalities of operations in order to discover similar Web resources.

The formal interface provides information to support the invocation of an operation. The formal interface is intended for machine consumption in order to facilitate automatic invocation. The SOAP-based Web Service is originally described with a formal interface. Hence, a client program can be automatically generated to invoke operations in a SOAP-based Web Service. In contrast, other Web resources (e.g., the HTTP-based APIs) do not have a formal interface and the SOA professionals need to manually write the request to invoke the operation. The fields contained in the formal interface depend on the type of Web resources. To invoke an HTTP-based API operation and a Web form, the URI, the HTTP verb, and the input parameters of the operation are required. The formal interface is described in an XML format which can be interpreted by a machine to automatically invoke the operation.

An excerpt is taken from the existing description of an operation. It provides more readable and detailed information, such as examples and demonstrations. It also offers a shortcut for a SOA professional to understand the operation without having to search for the operation in the entire document. The excerpt is available only for SOAP-based Web Services and HTTP-based APIs.

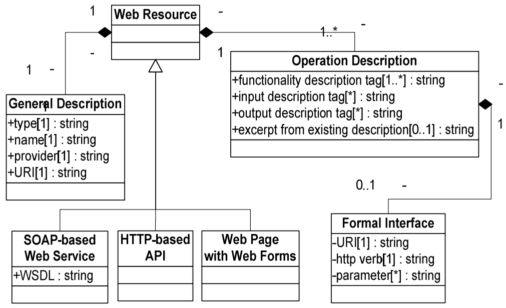

Figure 2: A unified resource description scheme

2.2 Constructing Resource Graph

A resource graph represents all resources using RDF model. It is a semantic network model, which consists of entities and relationships. Entities are identified globally with URIs. We want to

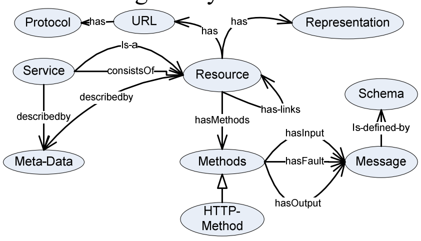

Figure 3:Conceptual model for RESTful Services

represent all the service in REST style. RESTful Service and RDF both represent "Resource," which is the main motivation behind using RDF. RDF provides a common framework for expressing information so it can be exchanged between applications without loss of meaning.

Figure 3 shows the conceptual model of the RESTful services. We model all SOAP based Web Services, HTTP-based API in terms of

RESTful services. As shown in figure 3, each service has one or more resource and service itself can be treated as resource. Each resource is uniquely identifiable and can have one or more than one representation. The information about the service is stored in the meta-data and resources have links to other resources. The input and output message has one or more parameter which is defined by the schema. Resource may link to other resources. We used link specification [7] to indicate the relationship between the resources. The “rel” attribute in the link specification give the information about the semantics of the link. In addition to types of “rel” defined by IANA [8], we introduced few other types that helps to represent resource in RESTful services more easily.

- **see-also** recommends another service
- **same–as** provides the similar services
- **is-a** defines is-a relation between the resources
- **contains** defines different service as in the case of composite service.
- **is-container-of** defines the resource-to-container relationship.

These relations help recommend services, identify similar services, and define the relationship between the resources. Semantic relationship helps to abstract the resource representation. Figure 4 shows the different resource and the semantic and data-link relations between the different resources. Since double bedroom and single bedroom has is-a relationship with room all the data link and semantic relation from room is carried over to those two different categories of the room. In addition to that, we added “method” attribute in the link. Thus, the user agent can know next possible resources to visit and along with method used to visit that resource. The information provided by the method attributes helps to define the flow. It also tells about the data required by another resource that end-user wants to visit. For example in Figure 4, “review” resource requires the information regarding the resource hotel. Hence, this type of relationship is called data-link relations. The small circle represents which method to can be invoked on resources from the current state. In Figure 4 the user can invoke only GET method in the resource review from the hotel resource. This substantially increases the user agents’ capability of discoverability of resource. In our resource graph, the resource from where end-users can start consuming service (a starting point) is defined as initial node. In Figure 4, hotel is represented in different color and it denotes the initial node. When a user requests a service this node is returned and from that node, user can start using the service.

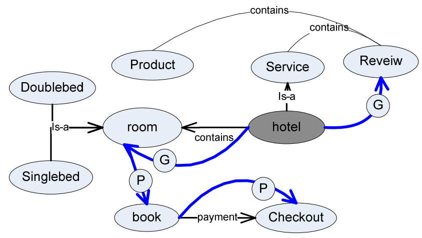

Figure 4: Resource and their relations

For example, if we want to book a room we need information regarding the resource room, Thus using REST approach it can be used as shown in following example:
As shown in Figure 5, when the resource room is requested, the output contains information about the next resource. The next state in this example can be one of the room types which are single bedroom and double bedroom. The semantic relation between these resources is give by “rel” attribute .After knowing the next resource user agent also needs to know which method to use. It is provided in method attribute in the link.

```
GET /hotel
HOST: foo.org
<various HTTP headers>
Output
<rooms>
<room id="singlebed" rate="40" currency="USD">
<link rel="isa" method=”GET” href="/rooms/singlebed/">
</room>
<room id="doublebed" rate="40" currency="USD">
<link rel="isa" method=”GET” href="/rooms/doublebed/">
</room>
</rooms>
```

Figure 5: Example of RESTful Service

WSDL/SOAP based services are lifted to resource level by identifying the resource and the HTTP method for each operation and then described in RESTful approach. When the end-user uses these resources the corresponding WSDL operations have to be invoked which can be identified using the mapping information for operation and resources. Converting every service in REST style helps the end users to see everything as resource and each resource is associated with the corresponding four methods. End-users are exposed to only the resource model of the service described in RDF. It is easier to deal with one model than dealing with different kinds of heterogeneous specification. Once the service is represented in the RDF form, it adds a lot of expressiveness with query languages (e.g., SPARQL [4]), transformation languages (e.g., GRDDL [5]), and rule languages (e.g., RIF [6]).

2.3 Inferring Ad-hoc Processes from Resource Graphs

Figure 6 illustrates the definition of ad-hoc processes. An ad-hoc process is characterized by a set of work items and sub ad-hoc processes performed by end-users to fulfill a goal. A work item in the ad-hoc process is a set of tasks which are collaborated together to accomplish a transaction. The work items in an ad-hoc process are connected through the relation defined by users or the semantic relations defined in the resource graph. A task is the combination of resource and operations. The resource in a task is defined in the resource graph. The operation in a task processes the resource. In RESTFul services, we can use operations Get, Post, Put and Delete. For example, the task "search for flight ticket" could be described as the resource "flight ticket" and the associated operation "get". Eventually, a task can be performed by one or several concrete Web services. Web services can be represented as resources in the resource graph. Therefore, we can trace the resource graph to find the associated Web services.

In our definition, work item is different from ad-hoc process although they both contain tasks. The resources related to the tasks in a work item are connected by data-link relations in the resource graph. The relations of tasks in a work item are very closely coupled. To fulfill a transaction, the user needs to execute all the tasks linked by the data. For example, the work item "buy flight tickets" includes tasks "search ticket", "choose ticket" and "pay the bill". In order to accomplish the goal of "buy flight tickets", the user has to perform all the three tasks in the work item. On the contrary, the relations of work items in an ad-hoc are loosely coupled. Different users can have different ad-hoc processes for the same goal. For instance, when planning a trip, some users may prefer taking flight than driving car, but other users may prefer driving instead of taking flight. The ad-hoc process of planning a trip for these two groups of users is different since the ad-hoc process for the former group contains the work item "buy flight tickets" and the latter does not have.

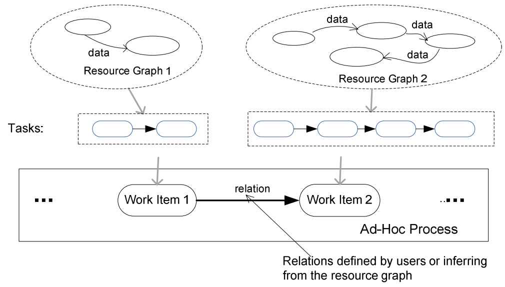

Figure 6: Definition of ad-hoc process

2.4 Inferring Work Item Relations from Resource Graphs

The work items in an ad-hoc process can be connected together using different relations. The relations of work items in ad-hoc processes are shown as follows.

- **Group** indicates that a set of work items should be performed together. The group relation can be further described as "And", "Or", "Sequence" or "Parallel" relations.
- **And** means that all the work items need to be performed. "And" relation does not provide the detailed information about the execution order. Thus the work items can be executed in any order by default. When we have more information about the execution order of work items, we can use "Sequence" and "Parallel" relations to describe "And" relations with the execution order information of the work items.
- **Sequence** means that the work items need to be executed following an order.
- **Parallel** means that the work items can be executed at the same time.
- **Or** indicates that the user only needs to execute one of the work items in the group.
- **Ungroup** releases the existing "group" relation. When a user does not like an existing group of work items, the user can use this relation to show that he/she does not agree to group these work items together.

Our framework analyzes the relations in resources graph and infers the work item relations from the resource graph. Table 1 shows the mapping from the resource graph to work item relations.

In Table 1, the "See_also" relation in resource graph can be use to recommend work items to users. The user has the option to perform it or ignore it. The “Same_as” relation in the resource graph indicates that two resources are equal. Therefore, we convert the “Same_as” relation into “Or” relation of work items. The siblings of “Is_a” relation in the resource graph are converted to “Or” relation since “Is_a” relation shows that one resource is an instance of another and these instances have the same features. The elements in a “Contains” relation in the resource graph is converted to “And” relation in the ad-hoc process.

Figure 7 gives an example to illustrate the main idea of inferring work item relations from the resource graph. In Figure 7, the resources “Hilton”, “Holiday_Inn”, “Flight”, and “Restaurant” are converted to work items in the ad-hoc process. In this example, if we trace entire resource graph, these work items can be implemented by several tasks which are associated with detailed resources.

Table 1: Infer work item relations from resource graphs

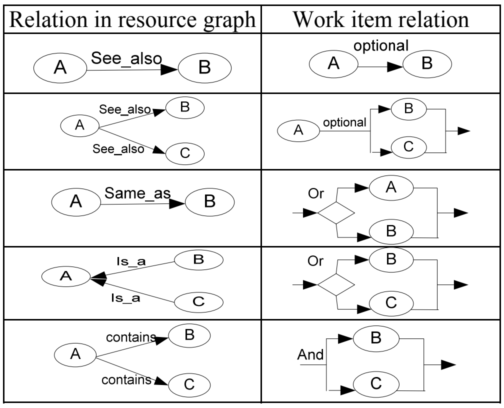

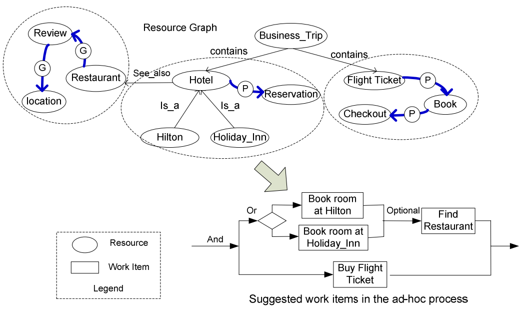

Figure 7: An example of relation inference

3 Conclusion

This paper presents a framework to compose heterogeneous resources on the Web. In our framework, Web resources can be described by the unified description schema and can be wrapped to RESTful services. The resources in our framework have semantic relationship and data link relations between them and can be described in RDF. RDF also adds a lot of expressiveness with query languages, transformation languages, and rule languages bringing more participation from end-users side. Thus we provide a framework whereby Internet end-users can compose their own Internet space, which is defined as a collection of resources that can be used to feed, filter, compose, disseminate, and reference information, data, and services to end-users according their profile, context, and mode of operation. By analyzing the relations among services, we can infer the ad-hoc processes to compose resources. The way of describing everything in terms of resources solves the integrating issues and drives the innovation in the end user side.

References

[1] D. Beckett, B. McBride (editors), “RDF/XML Syntax Specification (Revised),” W3C Recommendation (2004)
[2] R. Chinnici, J. Moreau, A. Ryman, S. Weerawarana, “Web Service Description Language,” W3C Recommendation (2007)
[3] Dublin Core Metadata Initiative, http://dublincore.org/, last accessed on October 12, 2010
[4] SPARQL Query Language for RDF, http://www.w3.org/TR/rdf-sparql-query/, last accessed on October 12, 2010
[5] Gleaning Resource Descriptions from Dialects of Languages (GRDDL), http://www.w3.org/2004/01/rdxh/spec, lat accessed on October 12, 2010
[6] RIF In RDF, http://www.w3.org/TR/rif-in-rdf/, last accessed on October 12, 2010
[7] Web Linking, http://tools.ietf.org/html/draft-nottingham-http-link-header-10#section-4, last accessed on October 12, 2010
[8] Link Relations, http://www.iana.org/assignments/link-relations/link-relations.xhtml, last accessed on October 12, 2010
[9] R. Fielding. “Architectural Styles and The Design of Network-based Software Architectures,” PhD thesis, University of California, Irvine (2000)
[10] Vinoski, S.: RESTful Web Services Development Checklist. Internet Computing,IEEE 12, 96–95 (2008)

Service Subscription for Personal Web

Chunyang Ye
University of Toronto
Toronto, Canada
chunyang.ye@utoronto.ca

Hans-Arno Jacobsen
University of Toronto
Toronto, Canada
jacobsen@eecg.toronto.edu

1. INTRODUCTION

Niklaus Wirth ever wrote down in his book a famous philosophy about computer programs, that is, *Algorithms + Data Structures = Programs* [7]. To date, the representation of data structures has evolved from functional languages, object-oriented programming, relational database, semi-structured data(XML) etc to *Semantic Web* [5]. Semantic Web is a web of data. By representing information as linked data and making them referable, Semantic Web provides a unique interface to access heterogeneous and distributed data sources from all over the world.

On the other hand, the implementation of algorithms has also been shifted from basic building blocks in functional languages to more well encapsulated object-oriented classes and components. Nowadays, the service-oriented architecture (SOA) is a new software engineering paradigm to design, implement and publish reusable software components as Web services [3]. By composing together suitable Web services from third-parties, service consumers can quickly develop distributed Internet applications across different organizations.

In this position paper, we extend Wirth's philosophy with a new interpretation about Internet applications based on the combination of Semantic Web and SOA technologies, that is, *Services + Semantic Web = Internet Applications*. Different from traditional programs where data structures and algorithms are built local to programs, the Semantic Web and SOA technologies allow the data and services to be shared among customers through standard interfaces (e.g., RDF [4] and WSDL [6]), respectively. Therefore, users (customers) are allowed to customize their own applications by consuming well selected services and data over the Internet.

This new vision enriches the development and usage patterns for Web applications. In particular, with respect to Personal Web applications, end-users (customers) can publish their personal information (e.g., comments on a movie, online shopping of particular products etc) as linked data on their personal web and share the data to their friends (e.g., the social network in Facebook [1]). By subscribing to the update of linked data of interest, end-users can get notified of their friends' interesting activities. On the other hand, end-users may not be restricted to get the notification of linked data update only. They may subscribe to some services to manipulate the updated linked data when the data is ready. For example, users may subscribe to the recommended movie information from their friends in the social network website. They may also subscribe to the online box office service to book the tickets based on their online calendar data. When a movie is recommended, the subscribed online box office service will be invoked to book the tickets for the users automatically.

The challenges of this new paradigm for Personal Web applications lie in how to glue the Semantic Web and SOA technologies to develop customized applications for customers. On the one hand, since the Semantic Web and SOA adopt different standards and protocols to share and exchange data, the interoperability between both technologies is a challenge. On the other hand, the data and services usually belong to different organizations that are beyond the control of customers, how to coordinate the data and services is another challenge.

In the rest of this paper, we present our solution to address the aforementioned challenges.

2. OVERVIEW OF METHODOLOGY

As mentioned in Section 1, the challenges of our proposal lie in the interoperability between Semantic Web and SOA technologies and the coordination of data and services that are not under the control of customers. In this section, we briefly introduce our solutions to address the challenges.

In order to interoperate between the Semantic Web and SOA, we extend our previous work [8] to allow services access the data from the Semantic Web based on event interfaces. An event interface declares what kinds of events a service wants to expose and what kinds of events it wants to subscribe. An event is defined as a state change. For example, if some linked data in the Semantic Web is updated, an event can be generated and propagated to the services that use the linked data. On being notified of the subscribed events, the services can update the copies of the linked data inside the services. On the other hand, if a service updates the copies of some linked data, then events will be raised and exposed. The exposed events are then used to update the corresponding linked data in the Semantic Web. Therefore, as illustrated in Fig. 1, we provide our Polaris framework [9] to bridge the gap between the linked data in Semantic Web and SOA services. By declaring the state changes related to

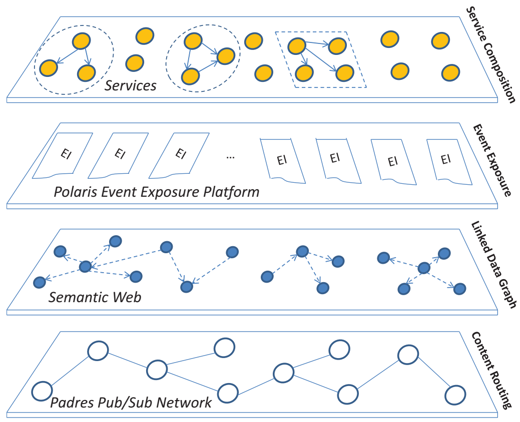

Figure 1: Overview of methodology.

the linked data of interest, services can access and synchronize with the linked data.

Another issue is to invoke the subscribed services when the subscribed linked data is ready. To do so, the Polaris event exposure platform allows users to design some Event-Condition-Action(ECA) rules. When an event is notified, the corresponding matched ECA rules are triggered to invoke the subscribed services transparently and asynchronously.

Since the linked data is distributed over the Web, an event related to the changes of linked data needs to propagated from the event producer to the event subscribers. Therefore, we use padres [2], a content-based pub/sub middleware to propagate events. Padres allows users to advertise and subscribe events based on their contents. When an event is published, it will be automatically delivered to all the subscribers whose subscriptions match the content of the event. With respect to Personal Web applications, users can subscribe to linked data from their friends by specifying the subscription with the social network information. As illustrated in Fig. 1, the Semantic Web built on top of the padres pub/sub network is allowed to route the changes of linked data to the subscribers transparently with social network information.

Therefore, by adopting the Padres pub/sub network and Polaris event exposure platforms to bridge the gaps between the Semantic Web and SOA services, our solution allow personal web users to customize their applications based on subscribed Semantic Web data and SOA services.

3. FUTURE WORK

The current work presented in this paper provides some preliminary functionality to prove the concept of service subscription for Personal Web applications based on our Polaris and Padres framework. In the future, we plan to complement the current work with the following features:

- **Service Subscription Language**. In our current work, users specify which services are invoked to handle the subscribed linked data using the ECA rules. This imposes difficulty for end-users to design customized applications. We plan to provide a service subscription language for end users to easily specify what services they need. The service subscription language should also be able to specify some non-functional requirements (e.g., a user may want to subscribe the services used by their friends). In this way, users do not have to find the services and link them to the Semantic Web data by themselves.

- **Service Discovery and Matching**. With given service subscription, we plan to develop solutions to discover and match the subscribed services. Different from the subscription of linked data, the service subscription may include the behavior description of services. This requires to extend traditional content-based matching algorithm to support behavioral matching. In addition, the matching algorithm may also need to rank and filter the services based on the social network information (e.g., candidate services recommended by friends are marked with higher priority).

- **Service Wrapper**. In order to make use of the linked data, services need to provide an event interface to describe how the data internal to a service and the linked data outside the service are exchanged. This task however is tedious and error-prone. Therefore, in our future work, we plan to provide a tool for services to generate wrappers automatically to map the data between services and the Semantic Web with event interfaces.

- **Linked Data Advertisement and Subscription**. In the Semantic Web, data from different locations is linked. The linking relationship between data forms linked data graphs. In some applications, if one data is changed, some other data linked to the changed data may also be needed. For example, if a user updates the linked data about his/her favorite movies, the information about the new movies are usually also needed by users. Therefore, we plan to explore the linked data graph to generate the related advertisements and subscriptions automatically. In this way, users can get the notification of related data transparently.

4. REFERENCES

[1] Facebook. http://www.facebook.com/.
[2] G. Li, V. Muthusamy, and H.-A. Jacobsen. A distributed service-oriented architecture for business process execution. *ACM Trans. Web*, 4(1):1–33, 2010.
[3] M. P. Papazoglou, P. Traverso, I. Ricerca, and S. Tecnologica. Service-oriented computing: State of the art and research challenges. *IEEE Computer*, 40:2007, 2007.
[4] W3C. Resource Description Framework. http://www.w3.org/RDF/.
[5] W3C. Semantic web. http://www.w3.org/RDF/FAQ.
[6] W3C. Web Service Description Language. http://www.w3.org/TR/wsdl.
[7] N. Wirth. *Algorithms + Data Structures = Programs*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1978.
[8] C. Ye and A. Jacobsen. Event exposure for web services: A grey-box approach to compose and evolve web services. pages 1–21. Springer-Verlag, LNCS Volume 6400, 2010.
[9] C. Ye and A. Jacobsen. Polaris: a framework to compose and evolve smart web services via event exposure. In *IBM CASCON Exhibits*, 2010.

MIDDLEWARE SYSTEMS RESEARCH GROUP MSRG.ORG

Service Subscription for Personal Web

Chunyang Ye and Hans-Arno Jacobsen
University of Toronto
2010

MIDDLEWARE SYSTEMS RESEARCH GROUP MSRG.ORG

Evolution of Computer Programs

Algorithms	+	Data Structures	=	Programs*
Functional blocks		Functional language		
Object-oriented classes		Object-oriented classes		
Software components		Relation database		
		Semi-structure (XML)		
Web Services		Semantic Web		Personal Web Applications

*N. Wirth. Algorithms + Data Structures = Programs. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1978

MIDDLEWARE SYSTEMS RESEARCH GROUP MSRG.ORG

Service-oriented Architecture

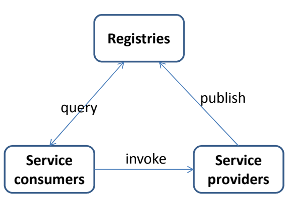

1) Service providers develop reusable software components and publish them as Web services in the public registries.

2) Service consumers query public registries, select suitable services and compose them to develop their own applications.

MIDDLEWARE SYSTEMS RESEARCH GROUP MSRG.ORG

Semantic Web

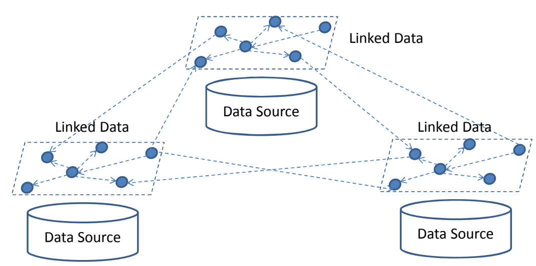

Semantic Web is a web of data, providing a unified interface to access data from different data sources over the Internet.

Personal Web

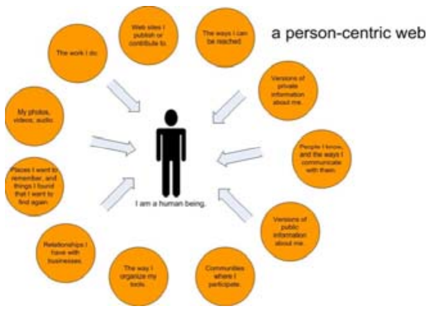

* This figure is from http://marynabadenhorst.globalteacher.org.au/pln/

5

Service Subscription

- How to customize services for Personal Web?
  - E.g., Subscribe to comments on the movies friends are interested in.
  - E.g., Subscribe to the online box office service to book the tickets based on the recommendations from friends.
- Challenges
  - How to combine Semantic Web and SOA?
    - Since they are beyond the control of users.

6

Overview of Methodology

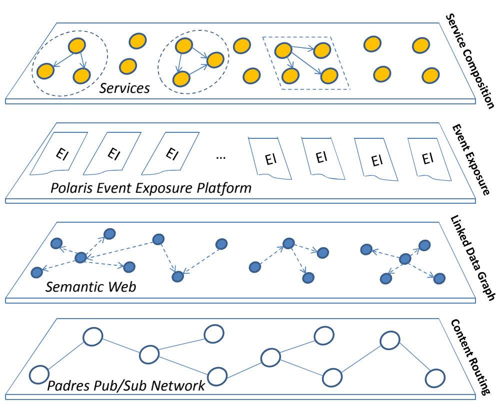

7

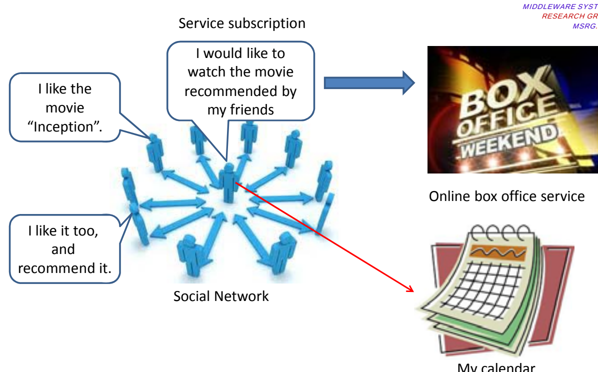

Application Scenario

8

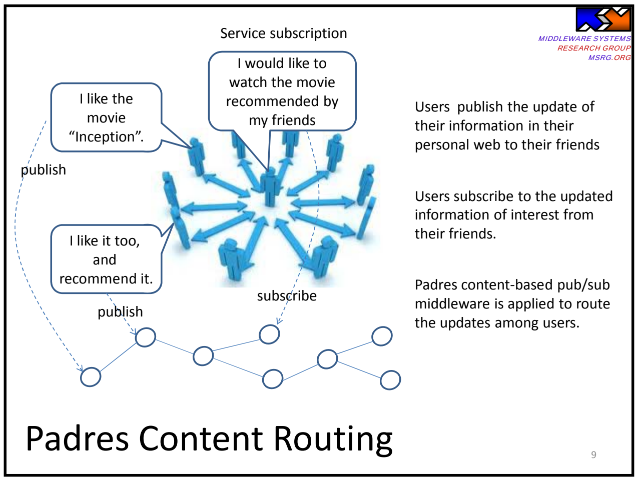
MIDDLEWARE SYSTEMS RESEARCH GROUP MSRG.ORG
Service subscription
I would like to watch the movie recommended by my friends
I like the movie "Inception".
publish
I like it too, and recommend it.
publish
subscribe
Users publish the update of their information in their personal web to their friends
Users subscribe to the updated information of interest from their friends.
Padres content-based pub/sub middleware is applied to route the updates among users.
Padres Content Routing
9

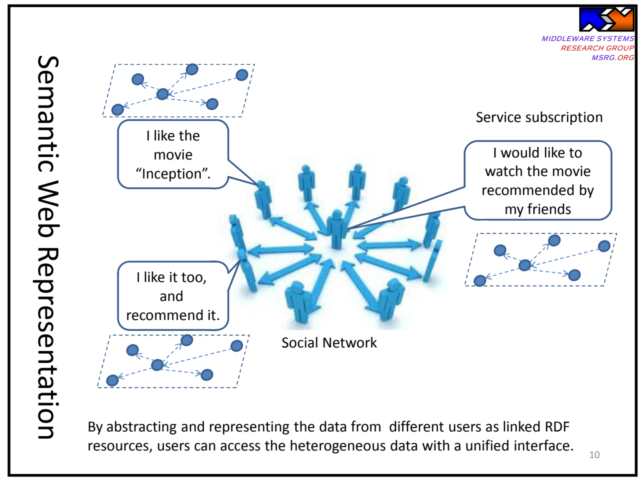
MIDDLEWARE SYSTEMS RESEARCH GROUP MSRG.ORG
Semantic Web Representation
I like the movie "Inception".
Service subscription
I would like to watch the movie recommended by my friends
I like it too, and recommend it.
Social Network
By abstracting and representing the data from different users as linked RDF resources, users can access the heterogeneous data with a unified interface.
10

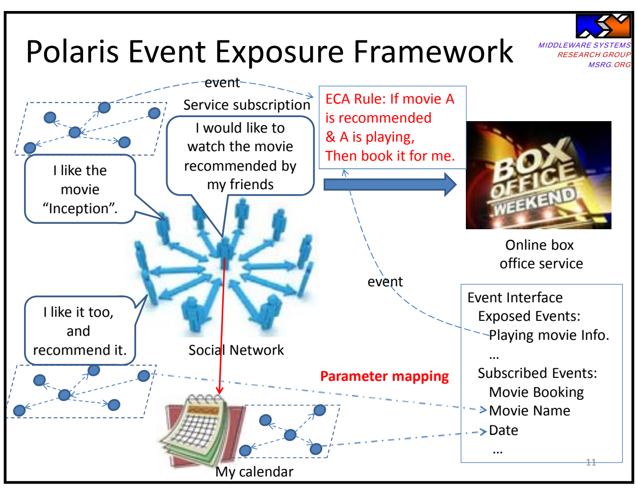
Polaris Event Exposure Framework
MIDDLEWARE SYSTEMS RESEARCH GROUP MSRG.ORG
event
Service subscription
I would like to watch the movie recommended by my friends
ECA Rule: If movie A is recommended & A is playing, Then book it for me.
I like the movie "Inception".
Online box office service
event
I like it too, and recommend it.
Social Network
Event Interface
Exposed Events:
Playing movie Info.
...
Subscribed Events:
Movie Booking
Movie Name
Date
...
Parameter mapping
My calendar
11

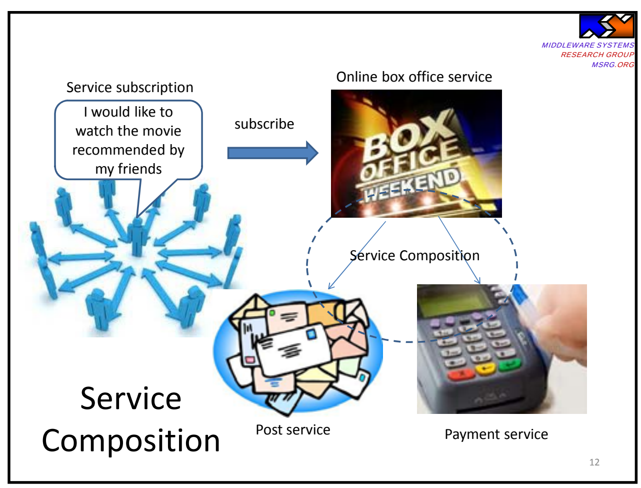
MIDDLEWARE SYSTEMS RESEARCH GROUP MSRG.ORG
Service subscription
I would like to watch the movie recommended by my friends
subscribe
Online box office service
Service Composition
Service Composition
Post service
Payment service
12

Future Work

- Service Subscription Language
  - Describe what services are needed
  - Behavioral perspective
- Service Discovery and Matching
  - Advertisement and Behavior Matching
- Service Wrapper
  - Automatically generate the event Interfaces to map the linked data and services' local data
- Linked Data Advertisement and Subscription
  - Correlation of linked data

Web Personalizer as User Consultant

Kamran Sartipi[†] and Mehran Najafi[♣]
DeGroote School of Business[†], Department of Computing and Software[♣]
McMaster University
Hamilton, ON, L8S 4M4, Canada
{sartipi, najafm}@mcmaster.ca

Introduction

An increase in the level of abstraction and ease of use in web-based computing requires more complexity and sophistication at the lower levels of implementation. Service oriented architecture (SOA) has provided more independency for the web service users, where the user can focus on the business logic aspects of the service integration. While the SOA technology has advanced the development of complex and distributed systems, it is still applied at the development phase and requires certain knowledge of system design and messaging techniques. The emergence of social network environments, user centric concept in information systems, and web application integration technologies such as Mashups, are driving forces for computer scientists and IT professionals to provide a customizable and effective environment for utilizing the rich variety of Internet resources. In this short paper, we propose different research avenues towards the overall requirements of the smart Personal Web systems. We propose the concept "Web Personalizer" which is a collection of three generic agents that are deployed at different platforms. These agents are specialized using roles and training skills to represent delegates from service provider to the user, or from user to the service provider to offer high-level services and consultations to the user.

The current state of the web applications is represented by the services that require extra knowledge and expertise from a normal user to take advantage of the available features and operations of the services. Given the large variety of web applications and the users' tight time schedules, the users will have to limit themselves to a minimum set of available service features. This is also the case in using other types of computerized systems such as automobile gadgets, home appliances and entertainment centers. In other words, the proper and efficient use of computerized systems (embedded or software based) requires an extended level of knowledge in different application domains. The user interaction capabilities of these systems are already sophisticated and hence these systems act as effective "user assistants" by providing different types of information to assist the users in performing their tasks. However, domain knowledge and expertise are still requirements for the user. The next generation of these computerized systems should be even more sophisticated by incorporating the required expertise as part of the system's functionality. This means a shift from "user-assistant" to "user-consultant". Therefore, instead of expecting the user of a SOA service to know significant amount of details and operational steps of using a specialized web application, the web service itself acts as an expert that consults with the user to provide an effective and customized use of the operations according to the user's specific context information. This provides an opportunity for the user to employ an expert (as Web Personalizer) to manage their web assets and perform desired tasks with minimum effort and time. Such a Web Personalizer provides smart interactions and ease of use for the user. The "Web Personalizer" agents are resident (as opposed to mobile agents that can move among different platforms) with customizable architecture that receive a set of well-defined task information in order to become expert and serve the user. The proposed Web Personalizers will be additions to the currently available services (i.e., traditional services), which receive a client's request for a service, perform the service at the provider's platform and return the results to the user.

Architecture of Web Personalizer

A Web Personalizer is a collection of three generic agents, namely "*service-agent",* "*broker-agent"* and "*analyzer-agent"* that are used in the context of an extended SOA architecture. Based on the traditional SOA reference model, service providers publish the descriptions of their services in the service registry. The service clients search the service registry to discover and connect to the select services. In this context, a service-agent represents a real-world delegate from the business owner to serve the client; a broker-agent and an analyzer-agent represent the client's delegates that communicate with the service provider on behalf of the client. Before being operational, these generic agents must be customized using three-part messages as <model, knowledge, data>. The model part is in the form of a business process or guideline that is intended to specialize the agent to act a specific role (e.g., financial advisor, sales person, home nurse, tester, auditor, or trace collector). The knowledge part is in the form of business rules and actions that provide the required expertise for the agent to perform the assigned role. The data part provides contextual information to be used by the business rules and actions. The structure of a generic agent consists of blocks such as "task manager", "business process engine", and "communication gates". The task manager receives the model part (hence it adopts a specific role) to control the sequence of the assigned business process. The business process engine engages the expertise by applying business rules and performing business actions. The communication gates are used to import the user's context data and to return the results of the desired operation to the client. Such a layered configuration for a generic agent provides ease of use and flexibility of configuration to perform different tasks for different users. It is assumed that different agents of the Web Personalizer have already been deployed at the intended platforms, i.e., service-agent, broker-agent, analyzer-agent in the client, broker, and provider platforms, respectively. In the followings, the steps for utilizing the Web Personalizer are presented.

Step 1: *Identifying the context of the user*. In service provisioning, *context* refers to any information that can be used to characterize the situation of a service requester or provider. We define a context as a tuple: <User, Role, User Location, Server Location, Time of Day, Team, Delegation, Requested Profile Status, Service Invocation Type, Requested Data Type, Login/Logout Event>. This context information is monitored dynamically to feed a database of context-logs which will be used during the service selection.

Step 2: *Selecting the required task*. The user asks for a specific task and the required expertise to assist her. Through mining of the context-logs (Step 1) and consulting with a web registry, a client proxy obtains a list of relevant services that provide different levels of expertise in that task, and generates a ranking lists of them with their capabilities and the associate charges. The user then selects an appropriate service, which best matches with their situation. In this context, the web registry should possess a list of application domains such as: banking, insurance, healthcare, telephone, airline, government, etc.; as well as a list of expertise within each domain, such as: mortgage consulting, car insurance negotiator, virtual nurse, TV technical support, on-line ticket reservator, PHR viewer, medication administrator, financial assistant, etc.

Step 3: *Delegating expertise to the client*. After selecting the required task, the client proxy retrieves the service descriptions of the selected service and invokes the service from the provider's platform. Instead of performing the requested task for the client, the provider will send a tuple of <model, knowledge, data> to the client where the generic service-agent will receive the tuple and customize itself (as discussed above) to become an expert consultant for the user to help them in an interactive mode of operation.

In the followings, the characteristics of the three agents of the Web Personalizer are discussed.

Expert Service Agent

The service-agent is a generic agent that is installed at the client's platform and is personalized as described in Step 3 above. It is the closest agent to the user in the proposed Web Personalizer and provides an expert service to the user. The service-agent ensures that the user will take advantage of the full capabilities of the available service functionality by adjusting its service locally and according to the user's context information. There are several advantages in processing a service locally as opposed to performing it at the service provider's platform: i) client data is confidential and revealing it to a service provider violates the client's privacy and security; ii) client data is usually large or changes over time, thus transferring it to a service provider increases the required network traffic; and iii) in the case of real-time or time-critical services, transferring client data to a service provider increase the response time. As an examples of a service-agent consider the case of a financial adviser in the context of stock market. To give proper advice, a financial adviser usually asks for personal information from their clients (e.g., client's portfolio or budget). By employing a service-agent to personalize a general advice for the client, the client does not have to reveal their personal information and hence their privacy is maintained [1]. Other example would be using a service-agent to assist a patient with a chronicle disease by following a flow-based guideline that has been prescribed by the physician.

Customizable Broker Agent

This agent is optional in the proposed Web Personalizer. It is meant to evaluate a set of candidate traditional web services in order to select for usage or aggregate with other services. The user will select from a list of high-level quality features, such as performance, security, or availability, which causes the required expertise be sent to the generic broker-agent. In this context, the web registry maintains a list of analysis expertise that is available for the user to choose from. The selected expertise will be sent to the generic broker-agent to customize it for the intended service evaluation operation. The broker-agent will then issue a number of service invocations to the candidate services that are provided by the user. The results will be sent back to the broker-agent, where it uses an objective function with parameters that are defined based on the client's contextual information. The broker-agent will then send back a list of ranked services to the user with a short report of the merits or drawbacks of each service. JaxView [4] provides a list of monitoring and security services that are based on monitoring the service traffic between the service client and service application. The broker-agent can enforce the business rules on the overall transactions between the client and server to maintain the client's integrity. This is done by ensuring the correct sequences of messages and checking the message contents.

Customizable Analyzer Agent

A generic analyzer-agent is intended to provide more in-depth analysis information for a customized broker-agent to perform sophisticated service quality analysis than those currently used for service selection and service aggregation. In this context, the broker-agent customizes the generic analyzer-agent, located at the service provider platform, to instrument the service application by embedding binary code into the service so that the analyzer can collect execution traces that belong to the broker-agent. The collected traces are returned to the broker-agent where different dynamic analyses can be performed on the execution traces such as security flaw identification or feature localization. For example, the analysis of patterns of execution traces using sequential pattern mining provides knowledge about feature scattering among the services. However, these analyses are intrusive and require the consent of the service providers. Such analyzers can also be specialized to act as policy monitoring, auditing, and testing agents by returning a variety of information from the SOAP messages. We are currently pursuing a research that uses analyzer-agents to perform dynamic analysis of web services based on pattern mining of the execution traces.

Domain Interoperability and Cross-Domain Interoperability

An important requirement for the proposed Personal Web infrastructure is the capability of domain and cross-domain interoperability. Such an infrastructure will allow the proposed Web Personalizers to integrate web services from legacy and new applications, that are deployed in heterogeneous platforms, and are located in the same or different application domains. Such interoperability requires a common information domain and a common semantic terminology system to allow different applications communicate and understand each other.

Domain Interoperability [2]. The migration of legacy web applications to a new architecture requires data and service reverse engineering to understand the structure and organization of the information system, comprehend the available standards for data and services representations, and perform the actual mapping. To tackle the complexity of network-centric interoperability, the trend is towards ease of use, vendor / language / platform independency, and in general raising the level of communication abstraction. The available technologies at different levels of abstraction such as: RPC/RMI (low-level), CORBA / DCOM (proprietary), XML and web services (non-proprietary), and enterprise level (SOA) to a large extent have solved the problem of interoperability of heterogeneous distributed systems. However, another challenge to be tackled is the interoperability of terms and concepts between different information organizations, namely semantic interoperability. This necessitates a different type of communication abstraction, which reduces the task of users from making frequent agreements on different concepts and services that they use locally. This issue can be resolved through standard terminology systems and services for exchanging specialized applications and workflows. Such interoperability issues have been extensively dealt with in the healthcare domain. Due to sensitivity of information in healthcare systems (mostly patient data) and a huge collection of terms and concepts, achieving interoperability among healthcare systems faces problems beyond those in most domains. Moreover, the growing cost of healthcare in most countries has been the driving force to put a lot of efforts in this domain for defining standardization processes. These efforts have resulted in developing international standards in common terminology systems (e.g., SNOMED and LOINC) and in organizing and representing the whole body of domain information into class diagrams (i.e., HL7 RIM), as well as a large set of standard message interactions to allow seamless communication among applications (i.e., HL7 v3 messaging). Therefore, any new attempt for interoperability among legacy healthcare systems should conform to these standards to ensure its compatibility and maintainability in the future. However, there are some problems in adopting HL7 v3 messaging standards due to its complexity and large message sizes. We have developed a prototype environment (MacSeie: McMaster Service-Based eHealth Integration Environment) for integrating two systems: a prototype electronic medical record system (EMR) that we developed, and an existing clinical decision support system (CDSS) which provides a collection of algorithms that use the patients medical history, current medications, allergies, and vascular risk and generates messages which contain recommendations for the patients and physicians. The MacSeie environment is web service based and HL7 v3 compliant.

Cross-Domain Interoperability [3]. Currently, most research activities are focused towards standardization and interoperability among information systems within the same domain. However, an emerging challenge is to address the exchange of information among heterogeneous applications in different domains, such as healthcare and insurance. Cross-domain interoperability

requires extensive and in-depth analysis of the information domains and semantic models of the corresponding domains in order to develop a new set of standards for communication.
The Healthcare Development Framework (HDF) specification is a product of the HL7 Modeling and Methodology work group, and defines the general methodology for developing messaging standards. The HDF is aimed at achieving interoperability within the healthcare domain using a common set of message elements. While provision of semantic interoperability within the same domain is a nontrivial research problem, it is even more challenging to provide interoperability among systems in two or more application domains. We have proposed a framework to achieve message-oriented semantic interoperability across global application domains such as healthcare and insurance to ensure the conformance of the resulting common standard with the local domain standards. The proposed approach is an extension to the HDF framework and is based on: comprehensive information and messaging model of RIM; hierarchical concept and terminology representation of SNOMED; and message development framework of HDF, all under well engineered standards of HL7 v3. The proposed approach provides the detailed description of three phases of the semantic interoperability framework consisting of "collaboration point analysis", "design harmonization", and "dynamic design" processes. This approach allows the domain and technical experts to analyze and engineer two application domain standards in order to extract business rules and message elements of the portions of the two domains which are subject to collaboration. A core standard information model has been adopted as a common information domain (CIM) where the information and concepts from two domains will be translated into the CIM. The most challenging part of the semantic interoperability provision is to identify and map between message elements of two XML message schemas with different structures. This mapping algorithm is interactive and ensures a precise mapping between two XML messages in terms of semantics of their tree nodes.

References:

[1]. A Framework for Context-Aware Services Using Service Customizer. M. Najafi and K. Sartipi. IEEE International Conference On Advanced Communication Technology. (ICACT 2010), Volume 2, pages 1339-1344. February 7-10, 2010, Phoenix Park, Korea.

[2]. Standard-based Data and Service Interoperability in eHealth Systems. K. Sartipi and M.H. Yarmand. IEEE International Conference on Software Maintenance (ICSM 2008), pages187 - 196. Beijing, China. September 2008.

[3]. Cross-Domain Information and Service Interoperability. K. Sartipi and A. Dehmoobad. ACM International Conference on Information Integration and Web-based Applications & Services (iiWAS 2008), pages 25-32. Nov 24-26, Linz, Austria.

[4]. JaxView: http://managedmethods.com/

WSCells for the Personal Web

Douglas Martin James R. Cordy

School of Computing, Queen's University, Kingston, Ontario

Abstract

In our previous work [1], we developed WSRD – a tool that analyzes WSDL service descriptions to gather related pieces of operations into self-contained units and put them into a standalone description we call a Web Service Cell, or WSCell (pronounced "wizzle"). WSCells are complete, self-contained single service operations that can be analyzed, compared and (in theory) used independently of their original services. In this paper we explore the role that WSCells may be able to play in the vision of the Personal Web.

1 Introduction

Web services described using the Web Service Description Language (WSDL) contain descriptions of multiple operations that are co-mingled throughout the service description. The interpretation of an individual operation depends on both its context in the description and the many remote pieces, such as types, portTypes and bindings, on which it depends. Isolating a particular operation description for analysis or service composition can be a tedious task.

In our previous work [1], we developed a tool we call WSRD (Web Service Restructuring of Descriptions) that collects and consolidates the information in a WSDL file related to each operation and packages them into concise self-contained descriptions that we call Web Service Cells, or WSCells (pronounced "wizzles") for short. WSCells provide a finer granularity view of services by separating operation descriptions from the description of the entire web service.

WSCells were originally designed to aid in the analysis of web service similarity and support web service tagging automation. However, in the context of the Personal Web, we now believe that they may have a more interesting role in aiding the selection and composition of personalized web services. WSRD can be used to separate WSDL descriptions into their constituent operation WSCells, and then users or advanced algorithms can pick and choose the operations that work best for any given situation, environment or current matter of concern, and compose them with each other to form a custom, personalized new service (Figure 1).

In this paper, we provide a brief background of our work on WSRD (Section 2), as well as some new ways that we envision WSCells being used to compose personal services (Section 3).

2 Background

In our previous work, our original plan was to modify WSDL descriptions to be better suited for processing using clone detection methods, with

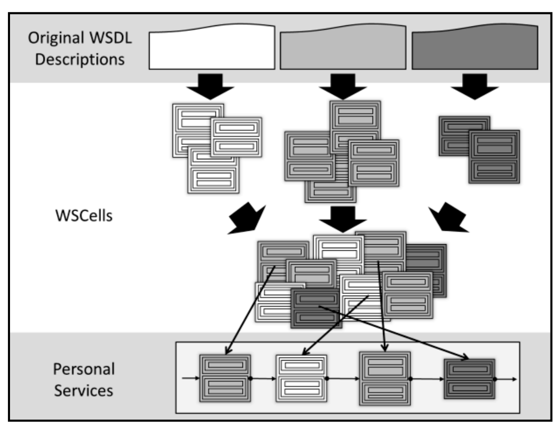

Figure 1
Separation of WSDL services into WSCells, and re-combination of WSCells into Personal Services

the goal of being able to identify and tag similar web service operations [1]. In this section, we briefly describe this technique using an illustrative example.

2.1 The Problem

Each WSDL file contains a description of one or more operations provided by a web service. Each of these descriptions is context-dependent and broken into several parts throughout the file, making it difficult to pick out and understand the entire description of a particular operation.

The complete description of an operation begins in the *portTypes* section of the WSDL file, where its inputs, outputs and possible faults are declared. Each of these refers to a *message* element, which in turn contains *parts* referring to elements in the *types* section. The types section contains an XML Schema that defines structured data types used for passing operation parameters. An element may have a type defined in the *types* section that contains elements of other types also defined there, and so on.

Thus the complete set of relevant pieces of the WSDL service description describing a particular operation are scattered throughout the WSDL file, and mixed in with irrelevant pieces of other operation descriptions. This makes it difficult to perform any sort of analysis of individual operations, or even of the service itself. This is where WSRD comes in.

2.2 WSRD

Using TXL [3], a source transformation language, we constructed a set of rules that transform the WSDL description of a web service into a set of self contained operation descriptions for the operations provided by the service. We call these self-contained operation descriptions Web Service Cells, or WSCells (pronounced “wizzles”) for short, because they are the like the cells that together form a web service. An example of one such WSCell is shown in Figure 2. Here we see an operation called *ReserveRoom* that takes a *Payment* object and a *Room* object as input and outputs an acknowledgement, with a fault in the case the room is unavailable.

WSRD includes a set of scripts that automates the generation of the complete set of WSCells for all WSDL files in a directory. Extracted WSCells can either be represented in a single file, or separated into individual files.

```
<operation name="ReserveRoom" >
  <input message="tns1:ReserveRoomRequest">
    <message name="ReserveRoomRequest">
       <part name="body"
             element="xsd1:ReserveRoomRequest">
         <element name="ReserveRoomRequest">
           <element name="payment" type="tns1:Payment">
             <element name="ccNumber" type="xsd:int"/>
             <element name="cardHolder" type="xsd:string"/>
             <element name="expiryDate" type="xsd:date"/>
             <element name="PIN" nillable="true" type="xsd:int"/>
           </element>
           <element name="room" type="tns1:Room">
             <element name="roomID" type="xsd:int"/>
             <element name="numBeds" type="xsd:int"/>
             <element name="isSmoking" type="xsd:boolean"/>
           </element>
         </element>
       </part>
    </message>
  </input>
  <output message="tns1:ReserveRoomResponse">
    <message name="ReserveRoomResponse">
       <part name="body" element="xsd1:ReserveRoomResponse">
         <element name="ReserveRoomResponse"/>
       </part>
    </message>
  </output>
  <fault message="tns1:RoomNotAvailableException">
    <message name="RoomNotAvailableException">
      <part name="body" element="xsd1:RoomNotAvailableException">
        <element name="RoomNotAvailableException"/>
      </part>
    </message>
  </fault>
</operation>
```

Figure 2
Sample Web Service Cell (WSCell) generated by WSRD

3 WSCells for the Personal Web

While the original intention of WSRD was to make clone detection of WSDL descriptions possible, WSCells can be useful in other applications as well. The self contained nature of WSCells make them great for analyzing, comparing and mix-and-matching operations, and for storing all information necessary to understand and use an operation separately from the entire service.

There are 3 main ways in which we envision WSCells being used in composing personal services: *imperative composition* is driven entirely by the user; *declarative composition* is driven by the user’s goals; and *self-driven composition* is driven entirely by the system as a cellular automaton. Each of these three possibilities is illustrated in Figure 3.

3.1 Imperative Composition

Because WSCells contain all the information needed to use a web service operation, they can be thought of as beads or blocks that can be piled or strung together like a child’s blocks. A user can

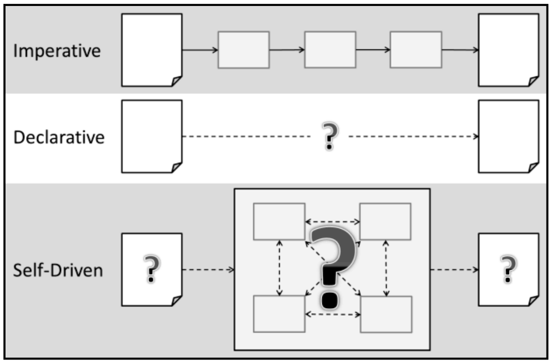

Figure 3
Three kinds of personal service composition

easily, given the right tools, pick and choose WSCells and hand compose them in a way that suits their particular needs to form a personalized service. Repetition, forking and joining structures can be imagined to form more complex combinations of operations. This is similar to the idea of "scripting" service composition in a higher-level notation.

As a concrete example, consider a generalization of Yahoo Pipes [2] where, instead of boxes that transform data streams like RSS or Atom, the boxes represent web service operations (WSCells) from a variety of web services. Users can connect the outputs of one operation to the inputs of another or, for a constant value, connect a text box with a value to an input. An example of this can be seen in Figure 4. In the figure, WSCells are represented using a combination of block notation from Yahoo Pipes and classes from a UML class diagram.

3.2 Declarative Composition

WSCells could also be automatically combined according to a user's goals, using a declarative approach to personal service composition that would iterate through the set of available WSCells looking for a path from the currently known information to the user's desired goal information. The implementation would be much like that of Prolog [4] – using unification and backtracking technology to "solve" the problem, without the necessity of the user explicitly specifying how to do so.

WSCells work well in this situation, because they give a concise, self-contained representation of the structure of a web service operation that can easily be matched or connected

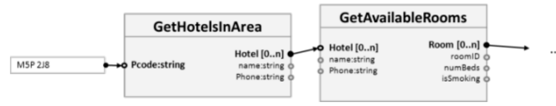

Figure 4
Example interface for imperative composition

to others. More specifically, they give a clear representation of the operation's inputs and outputs complete with their types and constraints. This makes it easy to match and explore whether or not two operations can be composed, in a way similar to the way Prolog unifies terms. All information about input and output for a particular operation is contained directly in its WSCell, meaning there is no need to search the entire WSDL description, and potential combinations can be discovered and explored efficiently.

3.3 Self-Driven Composition

While imperative and declarative composition rely on user input for composing services, self-driven composition relies on a cellular automaton model to discover new and interesting compositions to offer to the user. In this model, WSCells become active participants, looking for "partners" that they can combine with in meaningful ways to form new, higher-level "molecular" WSCells which can in turn look for others to combine with, and so on, gradually discovering the set of combined services that they can offer the user.

When combined with constraints to identify "interesting" combinations and reject uninteresting ones, this model of "emergent" service combination discovery could be interesting for the future.

4 Conclusion and Future Work

We have shown how our tool WSRD, or more specifically the self-contained WSCells that it produces, could prove to be a useful new concept in the discovery and composition of personalized services for the Personal Web. Their self-contained nature make WSCells ideal for selecting and combining operations from range of services, either imperatively by direct user selection, or automatically using either declarative or emergent programming to achieve user goals.

We have stated that WSCells contain all of the information that is needed to use a web service operation; however this is only partially true, and more work remains to be done if our idealistic vision is to be achieved. For example, it can be seen from Figure 2 that WSCells do not contain the address of the web service to which they belong. The reason is that this information was deemed unnecessary for the original purpose of finding similarities and was purposely left out. Future work will be put into packaging elements of the description that were ignored in our original work to make it possible to call the web service using only the operation's WSCell.

Acknowledgements

This research was made possible by an IBM CAS Fellowship and is partially supported by NSERC.

About the Authors

Douglas Martin is a second year M.Sc. student and James R. Cordy is a professor in the School of Computing at Queen's University. They can be reached by email at {doug,cordy}@cs.queensu.ca .

References

[1] D. Martin and J.R. Cordy, Chapter 12. "Towards Web Services Tagging by Similarity Detection" in M. Chignell et al. (eds.), *The Smart Internet: Current Research and Future Applications*, Lecture Notes in Computer Science 6400, Springer Verlag, October 2010, pp. 222-240.

[2] Yahoo Inc., "Pipes: Rewire the Web." Yahoo Pipes. *http://pipes.yahoo.com/pipes/.*

[3] J.R. Cordy, "The TXL Source Transformation Language", *Science of Computer Programming* 61,3 (August 2006), pp. 190-210.

[4] P. Deransart, L. Cervoni, and A. Ed-Dbali, *Prolog: the Standard: Reference Manual*, Springer Verlag 1996.