WPI-CS-TR-05-14 July 2005

Delivering Relevant and Useful Information with IMPACT

by

Paul J. Timmins
Craig E. Wills

Computer Science
Technical Report
Series

WORCESTER POLYTECHNIC INSTITUTE

Computer Science Department
100 Institute Road, Worcester, Massachusetts 01609-2280

Delivering Relevant and Useful Information with IMPACT

Paul J. Timmins and Craig E. Wills
Computer Science Department
Worcester Polytechnic Institute
Worcester, MA 01609
{ptimmins,cew}@cs.wpi.edu

Abstract

With today's nearly ubiquitous access to information of all types, it is difficult to discern what is important from what is not. Mobile and fixed location users alike are faced with a tradeoff: either check many information sources frequently or miss important information that may impact the user. In this work, we define a new approach to information management, one that uses the context of the user's planned activities to determine the priority and importance of information. By taking into account the user's planned activities, it is possible for the user to stay informed with minimal effort and disruption.

This paper describes our project to design and build a context-aware computing system, the Integrated Management of a Personal Augmented Calendar Tool (IMPACT) that uses a user's plan to prioritize, organize and deliver information. Its goal is to inform the user of any and all information that will impact the activities that a user intends to perform. To accomplish these tasks, the user's plan, which contains the user's current and future context, is built up from calendar data, to do lists, and preferences. For each activity in the plan, information sources are queried and assessed according to rules. Information is assessed for level of impact (severity) and urgency. The urgency is used to prioritize information, which is delivered to the user at times and via means designed to minimize disruption.

1 Introduction

Staying in-touch and informed with the world about us is one of the greatest benefits of the Internet. News, emails, traffic reports enable us to plan and react with relevant information at our fingertips. However, despite nearly ubiquitous access to almost any information, it remains difficult to stay informed due to the overwhelming volume and rate of change of this information. This problem is one of "information overload," having too much information available to us, thus rendering us unable to discern what is important and unable to focus.

To avoid information overload, people are forced to reduce the amount and rate of information they receive. Web users, for instance, reduce their information consumption by checking web sites at a less frequent rate. This approach leads to the user missing or receiving delayed information that may be important. Advanced users may subscribe to notification services, choosing sources that are likely to provide useful information, such as weather near their home. These services are useful if well-matched for relatively static information needs of a user, but fail when the information needs for a user change.

Current browsing and notification approaches are inefficient for a variety of reasons. Browsing requires the user to choose a source and seek out useful information, requiring the user to focus their attention on the browsing. Notifications interrupt the user based on when the information becomes available, not based on when the user needs the information. Aggregation can help address the problem of monitoring a large number of sources, but does not help the user decide what information actually is important to them. The root of the problem is that the user must sift through information to determine what is useful and what is not. Users are therefore faced with a tradeoff; spend more time filtering information from a number of sources or monitor a smaller number of sources and miss useful information.

Systems that rely on static rules and preferences, such as receiving traffic alerts for a particular city, cannot account for the fact that the interests and priorities of a user are dynamic. One key factor in determining the usefulness and urgency of information is whether it impacts an activity or task of the user. From traveling to meetings to research to social events, what the user does determines a large part of what the user is interested in. For instance, weather in another city and traffic along a highway may impact travel plans. Emails on a certain subject or from particular people may impact plans for meeting and social activities. The user's plan of activities can therefore be used to prioritize information based on how significantly it impacts the individual activities, reducing the time and attention required to filter through a large number of information sources.

Incorporating a user's dynamic interests and priorities into a system that manages information to provide timely and relevant information will result in more confidence by users that they are not missing critical information and allow them to focus on their task at hand. In addition, users must trust the system to deliver the relevant information without unnecessary disruptions. To handle this need, the concept of a user plan can be applied to another domain: minimizing the disruptiveness of interruptions. As notifications are less disruptive during activity transitions, most notifications can be delayed until the user is either in transition or has a break between activities.

Another application of the concept of a user plan is in organizing information for reference and quick access. Observing that a user is likely to request information that relates to an activity the user will perform, such as contact information and driving directions. Associating information directly to activities, allowing the user to quickly find information that relates to an activity, again reduces the time and attention required to find a particular piece of information.

This paper describes our project to design and build a context-aware computing system, the Integrated Management of a Personal Augmented Calendar Tool (IMPACT) that uses a user's plan to prioritize, organize and deliver information. Its goal is to inform the user of any and all information that will impact the activities that a user intends to perform. To accomplish these tasks, the user's plan, which contains the user's current and future context, is built up from static preferences and calendar-based data. For each activity in the plan, information sources are queried and assessed according to rules. Information is assessed for level of impact (severity) and urgency. The urgency is used to prioritize information, which is delivered to the user at times and via means designed to minimize disruption.

With this brief description of our approach, the remainder of the paper is organized as follows. In Section 2 an overview of related research is presented. Section 3 describes the concepts and approach behind the IMPACT project. Section 4 outlines a prototype implementation and describes the lessons learned followed by the next planned steps in Section 5. Section 6 discusses how a system such as IMPACT should be evaluated. Section 7 discusses future work and Section 8 summarizes the work thus far.

2 Related Research

The IMPACT project is related to a number of diverse research areas. An overview of relevant work for each of these areas is described in this section, and is organized in the following categories:

- Research into the nature and impact of disruptions in different environments emphasizes the importance of minimizing disruption in our daily lives.
- Mobile computing research attempts to simplify and improve access to information at any time.
- Web personalization and adaptive web research share similar goals as the IMPACT project, namely improving access to the "right" information, often by filtering and transforming web content.
- Context aware computing research is exploring how applications can use additional information, such as the user's location.
- Web extraction and Semantic Web research aims to take the WWW from blocks of HTML to rich, minable information.
- RSS enables aggregation of information across information sources, simplifying the user's ability to find and monitor information of importance.

2.1 Interruptions

Interruptions can cause significant disruptions. Research in [15] shows that a significant amount (15-20%) of work effort is consumed by handling interruptions. To address this issue, Horvitz and Apacible "present methods for inferring the cost of interrupting users based on multiple streams of events ... and data drawn from online calendars" [13]. Other work observes that interruptions

during physical transitions, such as standing up, are "better received" by the user [11] while attempts to measure the mental load of the user through heart rate and other physiological cues to deliver less disruptive notifications [6]. The IMPACT system expands upon these ideas by making use of the user plan to determine expected activity transitions, and therefore be able to deliver notifications when the user has completed one activity and has not yet begun another.

2.2 Mobile Computing

A goal of the IMPACT project is to present information in a way that is accessible and appropriate, independent of a user's physical location. Mobile devices, including phones and personal digital assistants, now provide nearly universal access to a wide range of content. They are an important tool used to communicate and access information. Yet, the availability of content for these devices is quite limited and difficult to use [8]. Previous work has looked at ways of improving the mobile browsing experience [4], and transforming web pages for improved access by mobile devices [17]. These approaches may be limited in how much they can improve the mobile web experience, as pointed out by [3]: "A literal translation from the Web to wireless is inadequate. Merely squeezing data onto small screens detracts from the user experience on mobile devices. By adaptively learning users' preferences, all users can have easy access to a vast amount of information at any time."

2.3 Adaptive Web

Additional research has argued that a one size fits all web site model, which can be adapted dynamically for desktop and mobile users, will not address the different needs of mobile users [2]. Instead, a web site personalizer is presented that "automatically adapts and personalizes a web site." They also note that users exhibit different goals and behaviors, from goal-oriented searching to browsing/surfing.

Other work explores how to personalize web content by learning user's interests and adapting the site to fit [1]. This idea is also known as an adaptive web, whereby sites are transformed "from today's inert collections of HTML pages and hyperlinks to intelligent, evolving entities" [22].

2.4 Context-Aware Computing

Context-aware computing research has focused on how additional information can be leveraged to enhance the computing experience [7]. One of the most likely sources of context is the users current physical location [16, 24], which can be used to enhance a users web experience through location-aware applications. Other work has explored a service that delivers notifications based on the location of the user [21]. An event planner that leverages contextual information to make more informed decisions is described in [23]. Context prediction attempts to use history and other mechanisms to predict future behavior or context of a user [20]. These types of research highlight the value of integrating diverse sources of information to enable new types of applications, applications that can adjust to the user's changing needs and interests.

2.5 Web Extraction

A significant portion of web information is not easily used in applications, as it is structured for human readability in HTML. This makes it difficult to deploy an system that attempts to not only aggregate, but filter and prioritize information. Research into general mechanisms for extracting data from web pages [12], as well as web services [10], has provided techniques for extracting information with meaning directly from the web.

The Semantic Web initiative takes a different approach. Instead of inferring meaning from pages, information should be provided by web sites already organized and structured according to its meaning. By associating not only structure, but also meaning, to web content, the Semantic Web provides the foundation for richer applications to be developed around the web [19].

2.6 RSS

RSS (Rich Site Summary or RDF Site Summary, or even Really Simple Syndication) defines an XML schema for delivering a summary of information available from a given web site along with hyperlinks to the actual content [18]. By providing a common format used across many web sites, RSS greatly simplifies the aggregation of web information, from HTML-based news pages to indexes of podcasts. This aggregation can provide a single source of information for the user to quickly get updated on information that is relevant to the user. Rather than solving these problems directly, it is envisioned that the IMPACT system would use RSS and Semantic Web technologies to interact with web and other information sources.

3 Approach

The goal of the IMPACT project is to improve a user's ability to monitor a wide range of information sources, and thus stay informed, while reducing the level of effort and disruption it would take using existing methods. This goal is accomplished in three ways. First, information is prioritized based on its level of impact to the user's planned activities. Second, information is organized according to the user's plan, reducing the effort required to find information that is related an activity. Third, delivery of information to the user takes into account the expected level of disruption to the user.

This section explores the approach taken by the IMPACT project, beginning with the determination of what activities the user intends to perform—the user plan. Next, information sources are queried for potentially impacting information, and assessed according to a set of rules. Reference information that is related to particular activities is then associated to each activity. Finally, potentially impacting information is prioritized and delivered to the user, minimizing unnecessary disruptions.

3.1 Building the User Plan

A user plan can be thought of as a user's future context. It provides a basis for determining what can impact the user, what information the user is likely to reference, and when interruptions will not significantly disrupt the user. To build a user plan, The IMPACT system combines the user's

calendar with external preferences, and other sources such as a to do list. These are discussed in the following sections.

3.2 Calendar

Many people use electronic calendars, often integrated with email applications, to record upcoming events. These events fall into a wide range of categories, and calendars are maintained to different level of detail. The user plan is built up from the two types of events that have direct user involvement: travel and appointments. Other types, such as holidays and vacation, are found in calendars but have less direct user involvement. While calendars may be not kept up to date by some, we expect that the capabilities of the IMPACT system will motivate users to keep their calendar more up-to-date and accurate.

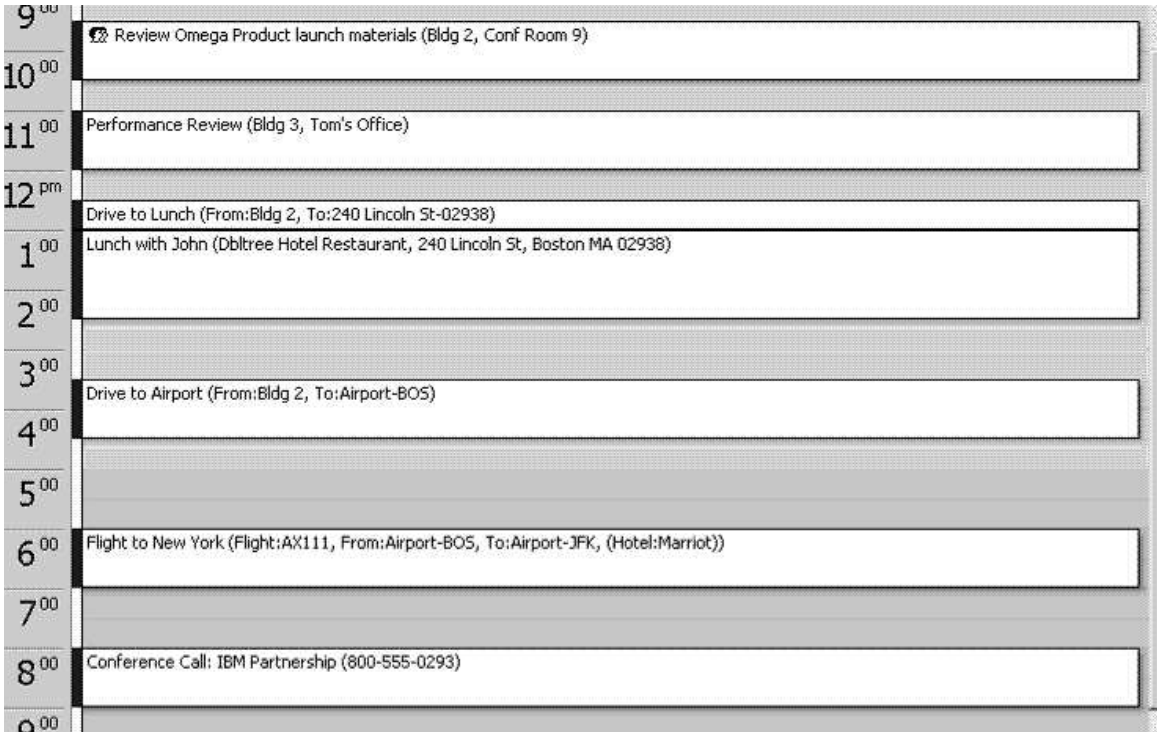

Figure 1: Typical User Calendar

A calendar provides a useful approximation of what the user intends to do throughout a day. Events of significance to the user are recorded in their calendar, including meetings, travel, and other appointments. Even if the user does not intend on attending an event that is contained in their calendar, information related to the event may be of relevance and interest. To illustrate the use of a calendar, Figure 1 shows a piece of a user's calendar with several types of appointments. In the morning, the user has meetings in two different buildings. The user then travels to lunch at a different location, and goes to the airport for travel to New York. The hotel the user will stay at is contained with the flight information.

3.2.1 To Do Lists

To Do lists provide a secondary source of unscheduled, but still planned, tasks that the user intends to perform at some point in time. Tasks in To Do lists have broader timeframes than scheduled activities, such as the user intending to complete a given task by the end of a week. Information that affects the tasks can impact the user and be of importance.

3.2.2 User Plan

The resulting user plan provides the schedule of events, and additional detail about them. Attributes from these events are augmented in the user plan with external information, such as the "home" address and the route from "home" to "work", providing a means of allowing reuse of common attributes. Table 1 illustrates the resulting user plan, from a portion of Figure 1.

0830	0850	Morning Commute
	Location Distance Route	Home to Work 20 miles Main Street to Rt 90 to St James Street
0930	1030	Review Omega product launch materials
	Location Attendee's	Bldg 2, Conf Room 9 Tom, Mary
1100	1200	Performance Review
	Location Attendee's	Building 3, Tom's Office Tom
TASK		Check on car loan
	Details Complete By	Tim R., (800)555-2222 Friday

Table 1: User Plan Example

Table 1 shows that the user commutes to work at 8:30am, with the first meeting of the day at from 9:30am to 10am. Two attendee's are invited, Tom and Mary. The check-mark shows the Tom has accepted, but Mary has declined the meeting. A "Performance Review" is scheduled at 11am. And, the user has a task "Check on Car Loan."

3.3 Information Aggregation

Once a user plan is available, related information is extracted for each event to provide a single, aggregated source of information for the user. The related information also provides the basis for determining whether and how to notify the user of specific information items. Web sites, desktop applications, and even file systems are types of information sources that can be related to particular events. Each type of information source requires a specific extractor, capable of retrieving information from the appropriate protocols and API. Source-specific rules define how the extractor should interact and extract data from a particular site. For example, to relate weather information from a web site a set of rules must be defined for an HTML-specific extractor. Extractors must be flexible

to support different types of relationships, from relating a cell in a table based on time of day to emails that contain a keyword. Information is then hierarchically associated with each event. For instance, weather would be a part of each event, with attributes of temperature, humidity, and type of weather. This approach provides an organized form for the extracted information.

Email, news, weather, traffic, and even files can all be related and relevant to a particular event. Recent emails from an attendee or similar to the subject are likely to be related to an event. Recent news or blog entries that contain keywords from the event, and even files can be related to the event. A flexible mechanism is needed to allow the user to define the relationships of interest for types of events, along with useful default rules. Search technologies can be applied here, or used, to process the rule and aid in the extraction of information.

Some information may not be fully extracted; rather a reference to the information may be maintained. A hyperlink or other type of reference would suffice, possibly with a summary of the information, in a manner similar to RSS.

3.4 Impact Assessment

With information associated to each event, the level of impact of the information is assessed next. A set of rules compare attributes of the event, the information, and possibly externally parameters for each information/event pair. This level of impact is a rating, from most impacting to least: critical, high, medium, low, none. These rules are based on the information itself, attributes of the event, or external parameters. Using weather again as an example: "if weather is heavy snow and event is travel by car, then impact is high."

The level of impact of information is dependent on a number of factors, dependent on aspects of the information itself and the activity. A travel activity, for instance, is impacted by information that changes the scheduling or duration of the activity. Bad weather, traffic, road construction, and the like can all affect the duration of driving between two locations. In addition, impacts can have a compounding effect. Delayed travel can impact a subsequent activity, requiring that the user either leave earlier or reschedule.

The concept of impacting information is not just a question of where, when and how long an activity takes place. It includes whether the information will affect how the user performs the activity. Press releases, research papers, blog entries, and news articles can all allow the user to perform an activity better. Newly available information can be directly related to some activity the user is performing, whether it is writing a paper, researching a topic, attending a meeting or making social plans.

3.5 Urgency Assessment

The urgency of informing the user is determined next, determining when the user should be notified. Urgency rates each piece of information as requiring immediate, delayed, or no notification. To make this assessment, the level of impact, proximity in time of the event, and the user's plan, and type of information are all taken into account. Immediate notification means that the user should be interrupted, regardless of the user's current activity. Delayed notification means that the user should be notified at the next convenient point in time. The rules are tunable, with users able to opt for fewer or more notifications, or suppress notifications of certain types of information.

3.6 Notifications

The delivery of the notifications can be performed through any means, configurable by the user and possibly dependent on the current user activity. Email, text messages, instant messages, and pop-ups all are appropriate depending on where and what the user is doing. A text message to a mobile phone is expected to be a common means, which for most phone services can be delivered via email. Periodically, the IMPACT system checks for pending notifications and schedules them based on urgency and prioritized by proximity in time.

Notifications may be bundled, delivered together, if they have the same urgency or relate to the same event. Therefore when an immediate notification is being sent, some delayed notifications may be sent with it. Multiple immediate notifications will be sent together. Bundled notifications minimize the overall disruptiveness of notifications by requiring fewer interruptions and deliver additional useful information with minimal added disruption. However, too much information in a single notification can increase the disruptiveness itself, thus the number of notifications in a bundle is limited to a constant, user-changeable parameter.

3.7 User Interface and Augmented Calendar

To provide access to the detailed user plan, augmented with related information, a web user interface is provided. The user interface provides an event summary and an impact summary. The event summary lists each event, related information for each event, and hyperlinks to further details. This organization facilitates quick access to information related to particular events. The second interface, the impact summary, organizes all impacting information. The impacting information is listed by event time and level of impact, enabling the user to quickly get up-to-date on related information. In addition, the user's calendar could be extended so that the related information can be access from within the calendar itself.

4 Prototype Implementation

To explore the core concepts of the IMPACT system, an initial prototype was built. Microsoft Outlook was used as the calendar source, with a few days of calendar information for sample data. The sample data was entered completely with all changes of location reflected by travel events so that the system did not have to infer travel. Each event was defined using conventions to simplify parsing. The subject of each event could contain several keywords for search purposes, in the form: <keyword:xxxx>. The location of each event could be defined as either a <street, city, state> or <street, zipcode>. For travel events, the source and destination locations were prefixed with the appropriate term, such as <source: street, city, state>, <destination: street, city state>. Air travel was encoded with the flight number <flight:AA111>, and source and destination airport code.

Calendar events were retrieved and exported in an XML structure using the Microsoft Outlook interop assembly. Only events that fell within the next 48 hours were retrieved. The retrieval was manually initiated for purposes of the prototype, but could be scheduled or triggered by changes to the calendar. The events were then loaded into event objects in the IMPACT system, a Java-based application with a Java Server Pages (JSP) front-end. Most of the raw attributes from the XML, such as location, required specific parsing and validation to convert to an appropriate data

structure.

Periodically, related information for each event was extracted from web sites. Extractors were written for weather, traffic from a state department of transportation web site, and air travel flight status. Each information type has a specific extractor that extracts text from an HTML DOM tree. This extractor works by finding exact and relative positions of text and keywords within a document. However, since a single web page often contains information that applies to multiple events, the content from a page was converted to an intermediate representation, such as weather for given zip code for the next few days, and put into a short-term cache.

As an example, the weather extractor retrieved data for a given zip code or city/state from the US National Oceanographic and Atmospheric Administration's web site, using specific query strings to pass the requested zip code. The resulting information, shown in Figure 2, was converted to an intermediate table that contained the weather for each time period for the zip code and could be reused later.

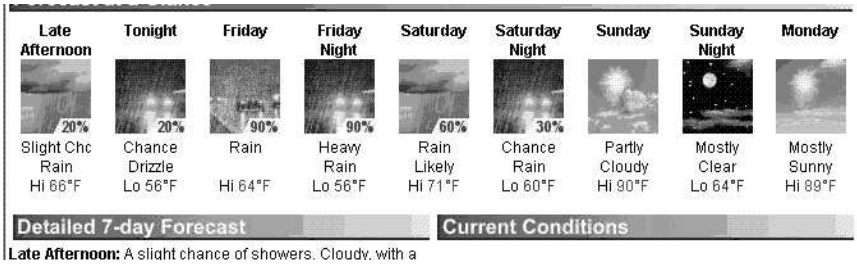

Figure 2: Weather from the National Oceanographic and Atmospheric Administration

The difficulty with this type of extraction, one which the Semantic Web initiative aims to address, is that the meanings of the text are not well defined. When exactly is late afternoon? When does tonight start? What are the allowable ranges of weather conditions? In addition, web pages are designed for human consumption, and thus the fact that the structure of most pages is generally stable is a side-effect of the design of particular pages, and not an inherent quality. The extractors written for the prototype were stable and accurate during short time periods, but changes to the web sites would likely require changes.

Built into the extractor was an assessment of the level of impact for each event that did not take into account the context or type of event. The level of impact was determined solely based on the content of the extracted information. For example, if the weather for the given time period included "heavy snow" or "alert," the level of impact was determined to be critical. A future system would also take into account other factors, such as traveling distance.

The web front end displayed upcoming events and related information. The related information was prefixed with the level of impact and highlighted with a color corresponding to the level of impact, with red text for critical and yellow for high impact. The notification system and urgency calculation was not present in this prototype.

The prototype showed the feasibility of the approach, that extracting and relating information to events could be done easily, assuming sufficient information was provided in the event. With a general HTML extractor, it required only moderate effort to add support for a new web site. The essence and contribution of the IMPACT project is the framework for relating information to particular events, and assessing the information in the context of the event. The simple rules

used in the prototype for determining level of impact provided a good first pass to demonstrate the concept, but would result in a large number of unnecessary interruptions.

5 Next Steps

To explore the benefits and effectiveness of the IMPACT system, further development is needed on the prototype to bring it to a stage suitable for user study and evaluation. Expanding the capabilities of the extractors, and developing extraction rules for commonly used information, is the immediate priority in order to have a broad set of information to relate to events and meet most user's needs. In addition, the prototype should be instrumented to facilitate a user study. This step would require tracking and logging capabilities, as well as a means for user input.

6 Evaluation

The objective of the IMPACT system is to improve the user's ability to stay informed with minimal disruption. Therefore, the evaluation of the approach should focus on its ability to deliver impacting information in terms of both completeness and accuracy. The IMPACT system should also provide a focal point for information that is likely to be accessed, by organizing information according to activities, and minimizing the user effort required to find information. In addition, the evaluation should assess whether the IMPACT system is effective at reducing the disruptiveness of notifications.

Evaluating the effectiveness of the IMPACT system in delivering impacting information begins with two key questions:

1. Is the determination of level of impact accurate? If the assessed level of impact differs from the user's assessment, then the system will not be able to properly prioritize notifications.

2. Is there impacting information missing from the system? If the system does not provide complete information then a user will be surprised and question the effectiveness of the system.

To measure these questions, user feedback is required in two areas. First, to identify missing impacting information, a user needs be able to point to information, such as by URL, and record the fact that the information impacted the user and if it is for a particular activity. An open question is whether impacting information can be related to an activity. Second, to determine whether the rule-based approach for determining the level of impact is effective, users need to be able to mark and correct inaccurate items. It may be that particular types of information require a more complex approach to assessing impact, for instance, the level of impact for weather may depend not only on the weather itself but the distance the user must travel.

Evaluating the effectiveness of information organized according to activities can only be partially measured. An expectation is that by making it easy to find information that is related to an event, that the user would rarely need to browse the same information sources directly. In other words, a user should not need to bypass the IMPACT system and browse the same information sources directly. In addition, it would be interesting to know how the IMPACT system changes a user's behavior. For instance, do users spend less time browsing and checking email?

Finally, evaluating the disruptiveness of notifications can be assessed by providing a feedback mechanism. First, user should be able to flag unnecessarily disruptive messages. Second, users should be able to record the disruptiveness of various notifications. The first type of feedback assesses whether the idea of allowing high impact notifications to be more disruptive than lesser impact notifications. The second type of information helps determine whether the user's plan is an accurate gauge of when notifications are more or less disruptive.

7 Future Research

While the IMPACT project focuses on information gathering and notification, other uses of future context may be used to reduce the frequency and disruptiveness of mobile computing devices.

Many problems arise when using a raw calendar, extracted from an application such as Microsoft Outlook, including inaccurate data, missing location information, and implicit travel (when adjacent appointments are in different locations). Any system relying on a user's calendar must overcome these issues to reside on a foundation of reliable information. To do so, the calendar may need to be maintained by external agents which retrieve information from other sources and update the calendar. For instance, air travel itineraries can be retrieved from a user's account at an airline web site and automatically inserted into the calendar.

Communication, whether through telephone, instant messaging or emails, can be highly disruptive to a current activity. Often, one would wish to suppress phone calls, except for emergencies, while in a meeting. There is no way to inform the caller of the user's current context or when they should call. Future context could be used here to inform other people of more appropriate times, such as letting a caller view what periods the callee is busy and free [5]. In addition, availability may be defined with current and future context in mind. For example, the user may choose to route personal calls into voicemail when attending business meetings, but always receive calls from the boss.

Another limitation, discussed earlier, is how to deal with unplanned activities and changes to planned events. The system ideally could identify an unplanned activity, and provide information that is potentially relevant. For example, if the user is driving along a major highway then traffic alerts along the highway are of interest. In addition, without knowing that the user is engaged in some unplanned activity the user may receive potentially disruptive interruptions. Unplanned activities can change the user plan unexpectedly and thus render it irrelevant or inaccurate. Research has found that unplanned activities make up a significant amount of our daily interactions [14]. These activities present a difficulty as there is no basis for relevant information to provide to the user nor does the system know when the user can be interrupted.

Furthermore, while the IMPACT system currently only provides information, the system could assist and aid the user in accomplishing an activity. With knowledge of the user's calendar and current context, the system could determine that the user does not have enough time to drive to the airport and make the flight, and automatically send a SMS message to the travel agent asking for help.

Determining what information is related to events through manual associations and rules may be a limited approach. Many relationships exist that are difficult to define and a user may never explicitly define them. It may be possible to learn some relationships, and infer the conditions in which information may impact an event. It might also be useful to share the rules and associations

across users to create a catalog of relationship rules.

A final area to consider is that of automatically managing and maintaining the calendar. Extracting and unifying other sources of event information, such as travel portals or even emails (for instance, "let's meet at 2pm @ Starbucks?"). Both of these are difficult information extraction challenges [9].

8 Summary

The IMPACT project seeks to aggregate and deliver information according to the user's plan, the activities the user is scheduled to perform. The project aims to reduce the effort required to monitor web sites and other information sources, minimizing interruptions while maximizing the information that is monitored.

We have built a prototype of the IMPACT system, exploring some of the underlying difficulties in such as system, including web information extraction and determining which information actually impacts the user. Further development of the prototype will be conducted, leading to a user study of the usefulness and effectiveness of the system.

References

[1] Massimiliano Albanese, Antonio Picariello, Carlo Sansone, and Lucio Sansone. Web personalization based on static information and dynamic user behavior. In *WIDM '04: Proceedings of the 6th annual ACM international workshop on Web information and data management*, pages 80–87, New York, NY, USA, 2004. ACM Press.

[2] Corin R. Anderson, Pedro Domingos, and Daniel S. Weld. Personalizing web sites for mobile users. In *WWW '01: Proceedings of the 10th international conference on World Wide Web*, pages 565–575, New York, NY, USA, 2001. ACM Press.

[3] Daniel Billsus, Clifford A. Brunk, Craig Evans, Brian Gladish, and Michael Pazzani. Adaptive interfaces for ubiquitous web access. *Commun. ACM*, 45(5):34–38, 2002.

[4] George Buchanan, Sarah Farrant, Matt Jones, Harold Thimbleby, Gary Marsden, and Michael Pazzani. Improving mobile internet usability. In *WWW '01: Proceedings of the 10th international conference on World Wide Web*, pages 673–680, New York, NY, USA, 2001. ACM Press.

[5] JJ Cadiz, Attila Narin, Gavin Jancke, Anoop Gupta, and Michael Boyle. Exploring pc-telephone convergence with the enhanced telephony prototype. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 215–222, New York, NY, USA, 2004. ACM Press.

[6] Daniel Chen and Roel Vertegaal. Using mental load for managing interruptions in physiologically attentive user interfaces. In *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*, pages 1513–1516, New York, NY, USA, 2004. ACM Press.

[7] Guanling Chen and David Kotz. A survey of context-aware mobile computing research. Technical Report TR2000-381, Dept. of Computer Science, Dartmouth College, November 2000.

[8] Luca Chittaro and Paolo Dal Cin. Evaluating interface design choices on wap phones: Navigation and selection. *Personal Ubiquitous Comput.*, 6(4):237–244, 2002.

[9] Oren Etzioni, Michael Cafarella, Doug Downey, Stanley Kok, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S. Weld, and Alexander Yates. Web-scale information extraction in knowitall: (preliminary results). In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 100–110, New York, NY, USA, 2004. ACM Press.

[10] James Hendler, Tim Berners-Lee, and Eric Miller. Integrating applications on the semantic web. In *Journal of the Institute of Electrical Engineers of Japan*, volume 122, pages 676–680, October 2002.

[11] Joyce Ho and Stephen S. Intille. Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 909–918, New York, NY, USA, 2005. ACM Press.

[12] Andrew Hogue and David Karger. Thresher: automating the unwrapping of semantic content from the world wide web. In *WWW '05: Proceedings of the 14th international conference on World Wide Web*, pages 86–95, New York, NY, USA, 2005. ACM Press.

[13] Eric Horvitz and Johnson Apacible. Learning and reasoning about interruption. In *ICMI '03: Proceedings of the 5th international conference on Multimodal interfaces*, pages 20–27, New York, NY, USA, 2003. ACM Press.

[14] Ellen A. Isaacs, John C. Tang, and Trevor Morris. Piazza: a desktop environment supporting impromptu and planned interactions. In *CSCW '96: Proceedings of the 1996 ACM conference on Computer supported cooperative work*, pages 315–324, New York, NY, USA, 1996. ACM Press.

[15] Thomas W. Jackson, Ray Dawson, and Darren Wilson. Understanding email interaction increases organizational productivity. *Commun. ACM*, 46(8):80–84, 2003.

[16] Tim Kindberg, John Barton, Jeff Morgan, Gene Becker, Debbie Caswell, Philippe Debaty, Gita Gopal, Marcos Frid, Venky Krishnan, Howard Morris, John Schettino, Bill Serra, and Mirjana Spasojevic. People, places, things: web presence for the real world. *Mob. Netw. Appl.*, 7(5):365–376, 2002.

[17] Björn Knutsson, Honghui Lu, Jeffrey Mogul, and Bryan Hopkins. Architecture and performance of server-directed transcoding. *ACM Trans. Inter. Tech.*, 3(4):392–424, 2003.

[18] Reuven M. Lerner. At the forge: syndication with rss. *Linux Journal*, 2004(126):8, 2004.

[19] Catherine C. Marshall and Frank M. Shipman. Which semantic web? In *HYPERTEXT '03: Proceedings of the fourteenth ACM conference on Hypertext and hypermedia*, pages 57–66, New York, NY, USA, 2003. ACM Press.

[20] Rene Mayrhofer. Context prediction based on context histories: Expected benefits, issues and current state-of-the-art. In T. Prante, B. Meyers, G. Fitzpatrick, and L. D. Harvel, editors, *Proceedings of the 1st International Workshop on Exploiting Context Histories in Smart Environments (ECHISE2005)*, May 2005. part of the Third International Conference on Pervasive Computing (PERVASIVE 2005).

[21] Jonathan P. Munson and Vineet K. Gupta. Location-based notification as a general-purpose service. In *WMC '02: Proceedings of the 2nd international workshop on Mobile commerce*, pages 40–44, New York, NY, USA, 2002. ACM Press.

[22] Mike Perkowitz and Oren Etzioni. Adaptive web sites: an AI challenge. In *IJCAI97: Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI'97)*, pages 16–23, Nagoya, Japan, August 23–29 1997.

[23] Zachary Pousman, Giovanni Iachello, Rachel Fithian, Jehan Moghazy, and John Stasko. Design iterations for a location-aware event planner. *Personal Ubiquitous Comput.*, 8(2):117–125, 2004.

[24] Bill N. Schilit, Anthony LaMarca, Gaetano Borriello, William G. Griswold, David McDonald, Edward Lazowska, Anand Balachandran, Jason Hong, and Vaughn Iverson. Challenge: ubiquitous location-aware computing and the "place lab" initiative. In *WMASH '03: Proceedings of the 1st ACM international workshop on Wireless mobile applications and services on WLAN hotspots*, pages 29–35, New York, NY, USA, 2003. ACM Press.