Automated Generation of Information Presentation for Pervasive Environments

Keith Houck Michelle Zhou
IBM T. J. Watson
30 Saw Mill River Rd. Rt 9A
Hawthorne, NY 10532
{khouck, mzhou}@us.ibm.com

ABSTRACT

One of the most challenging but largely unaddressed problems in pervasive computing is that of efficiently creating interactive information summaries that are individually tailored to a wide variety of pervasive devices and contexts. We are developing a new software entity to address this problem that will dynamically generate appropriate interactive information summaries that can be communicated by various devices in different contexts. Here we give a summary of our approach in the perspective of software engineering and stress several fundamental issues encountered in our research endeavor.

RESEARCH FOCUS & AN EXAMPLE SCENARIO

While other research in pervasive computing generally focuses on building hand-crafted information presentation for a specific environment [4], our work is unique in its emphasis on automating the creation of interactive information summaries tailored to a wide variety of pervasive devices. In particular, we employ pervasive devices individually or in a group to create a coordinated, interactive multimedia environment where users can take full advantage of their perception in information processing abilities.

Before describing our approach, we first present an example scenario to illustrate how our system will work.

Policeman Bob was on his routine patrol in his car, when the monitor near the front seat alerted him about a possible home invasion nearby. While the car navigator automatically displayed the route to the location, Bob asked for more information about the neighborhood (e.g., residents' information). Upon Bob's request, our system found the neighborhood information and summarized it using a geographical format (e.g., information overlaid on a map).

When approaching the destination, Bob noticed a suspicious vehicle parked in the driveway. To read the car license plate, Bob got out of his car with his radio-PDA. Nobody was in the car, and Bob used his radio to request a license plate check and indicated that the result be sent to his wrist watch. Bob input the plate number and also silenced his radio so as not to alert the suspect. By sending only visual alarms to his wrist watch, our system alerted Bob that the car was reportedly stolen by an extremely dangerous suspect. Since our system also knew that Bob sent the request using his radio, it notifies Bob that the suspect's picture can now be seen on his radio PDA. In addition, our system also knew that Bob has a desk-top computer in his office, so it then prepared a detailed multimedia information summary about the suspect and sent it to his office computer.

In the meantime, Bob examined the suspect's picture on his radio-PDA and used his wrist watch to request backups silently. Since Bob and his backups could not find the suspect nearby, he went back to the police station. To finish his paper work, Bob reviewed the multimedia summaries about the suspect sent earlier by our system.

RESEARCH CHALLENGES

As illustrated by this hypothetical scenario, our work focuses on creating an integrated information environment that involves two-way interaction between one or more users and a wide variety of devices in a dynamically changing context (e.g., from inside the police car to the suspected area). Unlike a conventional computing environment, a pervasive environment poses extra challenges:

First, it is a difficult and time-consuming task to handcraft information summaries for every different type of device because of the diversity of the devices involved. For example, Bob used his wrist watch, radio, and desk-top computer to gather information; each will be given a different information summary that meets the specific needs of the device. Moreover, multiple devices may be used together in a single scenario (the wrist watch and the radio), and some devices are very "minimal."

Second, it is also very difficult to tailor the information summaries and the interaction formats to pervasive contexts. Recall that we can only provide Bob with silent (visual) alert signals because of a muted audio channel. But back in his office, Bob utilized a full version of the multimedia summary, containing speech, audio, and graphics. Moreover, such context is dynamically changing as the users may often need to shift from one focus (e.g., location or tasks) to another.

OUR APPROACH

To create an integrated and pervasive information environment where users can process information in an effective and efficient manner, we are developing an intelligent information system. Our system, shown in Figure 1, has knowledge about the information to be presented, the presentation context, and the characteristics of various devices. Using its knowledge, it can automatically create tailored multimedia information summaries on the fly for a wide variety of devices. Automatically generating information summaries not only alleviates design burdens for application developers, but also allows us to incorporate real-time knowledge of the pervasive environment we are targeting. To accommo-

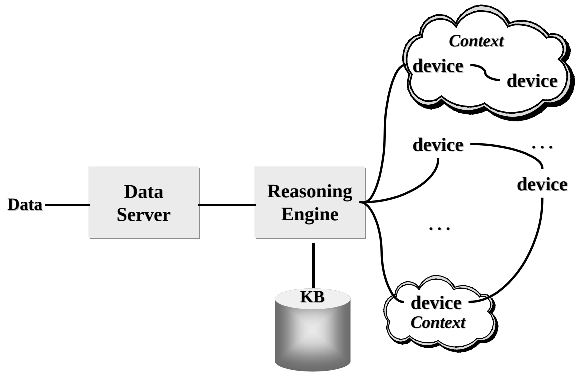

Figure 1. A simplified system overview

date the dynamic pervasive context, the system also has the ability to modify its planned multimedia summaries so the summaries can meet the new situations. Furthermore, it also coordinates the activities of multiple devices to let them work smoothly in concert. Here we address three fundamental issues that arise in our development process.

Data Modeling

To facilitate an automated generation process, we must provide our system with adequate knowledge of the pervasive environment, including the characteristics of various pervasive devices and contexts. So far we have modeled various devices along five dimensions to address their computational, communication, presentation, interaction, and storage capabilities (Table 1). All these five dimensions can be used as constraints to influence our system's decision as to what proper presentation or interaction metaphors should be used for a specific device. For example, if a device has very limited graphics capability, our system may not formulate an information summary using very sophisticated visual metaphors such as 3D animations.

Similarly, we also characterize the surrounding context including the user identity and tasks. As modeling and capturing context is an active research area itself [2], we don't discuss it further here.

System Intelligence

To cope with the diverse and dynamic nature of pervasive environments, we are experimenting with an intelligent software paradigm that can automatically generate interactive multimedia information summaries for a wide variety of devices. These information summaries are also tailored to specific pervasive contexts where users could fully take advantage of their perception as well as the capabilities of the devices. As in our example scenario, only simple visual

Computation	*Communication*	*Presentation*	*Interaction*	*Storage*
cpu	capacity	display	type	size
	speed	size	method	
		resolution		
		color		
		graphics		
		geoEngine		
		audio		
		channel		
		speaker		

Table 1. Device characterization

messages are created to alert Bob of a dangerous suspect, while a full multimedia summary of the suspect, including coordinated speech, text, and graphics, is created for Bob in his private office.

As a continuation of our previous work on automated visual presentation generation [6], we are extending our design engine to support a pervasive environment. However, the pervasive environment adds a number of new challenges to the automated generation problem.

Device and Context Adaptation. Unlike our previous reasoning engine that has a relatively simple model of the device and context (i.e., desktop computing), the reasoning engine must now take into account the constraints that arise from the capabilities of the pervasive devices and their application contexts. Based on our characterization of various devices, we abstract the correlations between the device characteristics and various presentation and interaction metaphors. For example, for simple devices (e.g., a wristwatch) that have very limited graphics presentation capability, we should use complimentary presentation metaphors (e.g., audio) or create collective presentations (see below).

In addition, the information summaries must also be tailored to a specific pervasive context. Using our context model, we also abstract general design principles that help select proper presentation and interaction metaphors for different contexts. For example, one of the principles may state that no audio input or output may be used in information summaries to be presented in a casual public setting (e.g., a train station).

Reactive Reasoning. In a conventional desktop environment, we usually assume that the surrounding context remains unchanged during the process of information summary generation and presentation. However, this is no longer the case in a pervasive environment because of the mobility of the pervasive devices. For example, users may request information summaries from a private setting (generation context), but may latter move to a crowded public setting and must view information summaries in that context. Moreover, unlike the stability we have enjoyed in a conventional computing environment, the unreliability of network connectivity or the device functions themselves also increase the unpredictability of a pervasive environment.

In these cases, the system must be able to perform reasoning under uncertainty and react accordingly [3]. In other words, the system may find that previous planned information summaries are inadequate for the new situation. To accommodate the new context, the system must modify the current plan or abandon it and start over. Starting from scratch could be very costly. Thus most systems always attempt to patch up the current plan to satisfy the new conditions. To patch up a plan, we must equip the system with sufficient knowledge of the presentation dialogue history: when the situation has changed and whether certain information has already been presented or not.

Device Coordination. One of the unique advantages of pervasive computing is that different devices could be used collectively to achieve one or more goals together. To conserve power on mobile devices (e.g., a cell phone) and compensate for their limited presentation capabilities, for example, we may use other devices (e.g., a laptop) to act as secondary information presenters. In other words, multiple information summaries may need to be created for many different devices at one time. In our police scenario, for example, several information summaries of the suspect's vehicle are created: one for the wristwatch, the radio, and the desktop at the police station.

To enable a synchronized view of multiple information summaries on various devices, we are investigating two coordination paradigms: automatic and manual synchronization. *Automatic synchronization* updates information summaries on all devices if there is an update on any of the devices. In our example scenario, if Bob uses his wrist watch to view the suspect information, the information summary created on his radio is also updated so that he could easily recognize the context when switching to his radio. In contrast, *manual synchronization* only updates an information summary on a device when explicitly asked to. For example, when back at the station, Bob may connect his wrist watch with his desktop computer and explicitly ask to update the desktop information summary to match the information currently displayed on his watch.

In both paradigms, the system must be able to compare different information presentation segments. Fortunately, our system can label different presentation segments in a summary by their content, since information summaries are created from scratch based on data semantics. Using these semantic labels, different presentation segments can then be matched and synchronized.

Performance Optimization

Inherent in our design are some fundamental performance costs. There are two major areas of concern:

- the time required for performing reasoning, and
- the extra communication flows incurred between the reasoning engine and various devices.

Fortunately, there are also some inherent opportunities for optimization that will help to mitigate the noted problems.

Leveraging our semantic knowledge of the data, we can predict what other information details are likely to be requested next and compute those in advance. Then, leveraging our knowledge of device capabilities, we can push these results as far down the pipeline (toward the target device) as allowed by available storage. In our police scenario, when Bob requests the information about the suspicious vehicle, the system can anticipate the follow-on request for more detailed information about the owner (e.g., criminal record) and automatically begin the process of querying, formatting, and forwarding this information. Since we have a deep understanding of the underlying information, cost-benefit estimates can be made to decide whether an anticipated query should be processed in advance at all.

By far the most difficult optimization is to distribute parts of the reasoning process onto the pervasive devices themselves. In addition to the computational overhead and possible power consumption issues for mobile devices, we are also concerned with the difficulty in making globally optimized decisions without having a global state. For example, independent reasoning processes on different devices may result in uncoordinated information summaries. Nevertheless, there do appear to be some cases where distributed reasoning is both feasible and effective. Consider the case of sending an asynchronous message to a cell phone device (e.g., Short Message Service). If the display of the cell phone is in the user's pocket, the cell phone would then need to find an alternative way of notifying the user. With a small amount of replanning in the phone, this message can be redirected to the user's other devices (e.g., wrist watch). Local replanning avoids the delay of sending the message all the way back to the original reasoning engine for replanning.

To improve the performance of the reasoning engine itself, we have employed an intelligent object-oriented approach. This approach has enabled local algorithmic optimizations to be implemented.

CURRENT STATUS & FUTURE WORK

We have a complete working version of a general-purpose information summary generation engine that utilizes a practical AI planning algorithm [5]. We have also implemented an initial version of a semantics-based presentation coordination component [1]. Using our device and context characterization, we are extending our engine to accommodate various pervasive device and context constraints. We are working with devices including wrist watches, palm displays, and hand-held displays (CLIO). We are also researching reactive planning methods so that planned information summaries can be modified to suit new situations. In the future, we would look into a distributed planning paradigm where multiple information summaries may be created simultaneously.

REFERENCES

[1] M. Dalal, S. Feiner, K. McKeown, S. Pan, M. Zhou, T. Hoellerer, J. Shaw, Y. Feng, and J. Fromer. Negotiation for automated generation of temporal multimedia presentations. In *Proc. ACM Multimedia '96*, pages 55–64, Boston, MA, November 18-22 1996.

[2] D. Salber, A. Dey, and Abowd, G. The context toolkit: Aiding the development of context-enabled applications. In *CHI '99*, pages 434–441, 1999.

[3] D. Wilkins, K. Myers, J. Lowrance, and L. Wesley. Planning and reacting in uncertain and dynamic environments. *Journal of Experimental and Theoretical AI*, 6:197–227, 1994.

[4] M. Wiser. Some computer science issues in ubiquitous computing. *Communications of the ACM*, 36(7):75–84, 1993.

[5] M. Zhou. Visual planning: A practical approach to automated visual presentation. In *Proc. IJCAI '99*, Stockholm, Sweden, August 1999.

[6] M. Zhou and S. Feiner. Automated visual presentation: From heterogeneous information to coherent visual discourse. *J. of Intelligent Information Systems*, 1998.