A MODEL FOR PERSONALISED COMMUNICATIONS CONTROL IN PERVASIVE SYSTEMS

Elizabeth Papadopoulou, Sarah McBurney, M. Howard Williams
Heriot-Watt University
Riccarton, Edinburgh,
UK
{ceeep1, ceesmm1, mhw}@macs.hw.ac.uk

ABSTRACT
As communication systems become more and more complex and integrate an increasing variety of different types of systems and devices, an important aim of pervasive systems is to provide users with more control over the delivery of different forms of communications. Simple forms of user control already exist in some systems, and work is ongoing to develop more sophisticated systems. Daidalos and Persist are two European research projects aimed at developing pervasive systems. In the process one aspect that has been investigated is the user control over communications in a general context-aware pervasive system. This paper outlines a general model developed to describe the processes, which was used as the basis of the implemented system in Daidalos. This model is being studied further in the Persist project.

KEY WORDS
Communications control, Pervasive, User preferences, Personalisation.

1. Introduction

The continuing rapid developments of communication technologies are enabling people to interact with each other and with a fast expanding set of data sources (including text, images, video streams, MP3 files, etc.) via a growing number of different devices at different locations. These developments are steadily moving us closer to achieving the goal of universal access – i.e. enabling users to have access to communications and data in accordance with their needs no matter where they are. Thus in the future users should be able to access communications ubiquitously through a variety of networks and stationary or mobile devices [1, 2].

However, achieving the goal of universal access to communications and data is only part of what is required of a communication system. It is also increasingly important to enable users to have more control over what information is delivered, how it is presented to them, on what devices it is presented (their own or ones located in their environment), when and how communications are delivered, and so on. Thus part of the vision of future communication lies in a user-oriented universal communication system that can accommodate versatile communication needs [3]-[5]. Such a system should not only be able to deliver information at any time, in any place and in any form (within reason) but also give users control over when, where and how communications are delivered. This must be achieved in a way that takes account of the user's preferences and context.

Another development is the proliferation of devices in the user's environment - including not only conventional computing and communication devices, but also a rapidly growing collection of computer embedded devices to monitor sensors, control tasks and access information sources. As the cost of these devices continues to fall and they become more powerful, there is an increasing demand for them to communicate with each other and access and provide information ubiquitously. This further reinforces the need for users to be able to control these communication flows as they affect him/her.

Thus within the general concept of pervasive computing [6], the need for control over communications becomes increasingly important. One aspect of the vision of a pervasive environment is the personalized control over what communications are delivered where and when. This involves more flexible routing of both data and communications (including the delivery of messages to a recipient independent of his/her location) and the allocation of data streams to appropriate devices to suit the user, with appropriate content adaptation where necessary.

The idea of personalised communications control has been investigated in some detail in the European research project Daidalos [7], and incorporated into a prototype pervasive system which has been used to demonstrate it. This paper presents a model for personalised communications control which can be used to guide implementation. The next section provides some background to the model, describing some scenarios and requirements derived from them. Section 3 presents the basic model and section 4 describes some extensions to this. Section 5 discusses some simplifications that can be made while section 6 concludes.

2. Background

2.1 Messages

Before looking at scenarios and requirements for personalised communications control, we begin by defining what we mean by a *message*. This is the unit of communication, which lies at the heart of the model. A message may be a traditional form of message (email, SMS, MMS, etc.), a voice phone call or even a videoconference or broadcast video stream. A message has a sender and one or more recipients. It transfers content in one direction or two (in the case of the voice call or videoconference).

The focus of this paper is a model for the personalised control over the delivery of messages. The notion of delivery of a message is interpreted as displaying it on an appropriate device (in the case of email, SMS, MMS) or connecting it to an appropriate device (in the case of a voice call or video stream).

2.2 Scenarios

This section describes three scenarios, which illustrate the types of personalised control of messages that are foreseen in pervasive systems.

(1) Anna the diabetic
Anna is a diabetic. To help her she wears a body sensor, which measures her blood sugar level (BSL) at regular intervals. If the value of the blood sugar level reaches a certain level (i.e. a medical emergency is detected – hyperglycaemic or hypoglycaemic), an alert message will be sent to her. She prefers this to be delivered on her PDA. If she does not respond to this message within a specified time interval, she has set up the system to determine if her parents or any of her friends are located in her vicinity and try to contact them with a slightly different message. If this fails it will notify the hospital ambulance management centre of her location.

The personalised message control here is triggered by the arrival of a message from the body sensor and decides who should be notified with what message and under what conditions (e.g. proximity to Anna).

(2) Bart the chauffeur
Bart is a chauffeur. While sitting at home one morning watching a newscast on his PC, he receives a phone call from his boss Susan. Since she is his boss, the system puts the newscast on hold and connects the call to his mobile phone. An important client, Jane, has arranged to fly in that morning and Susan asks Bart to collect her from the airport. Her flight is scheduled to arrive at 9.15 and he will need to set off at once to get there in time.

Still talking to Susan, Bart goes to his car. On entering the car, the voice call is transferred automatically to the car phone. Bart drives off while still talking to Susan. When the call is finished, the system transfers the suspended newscast to the car. Here it continues from the point it was suspended, except that it is directed only to the car speakers (audio mode) since Bart is driving. And so on.

(3) Carol the doctor
Carol is a general practitioner. She runs her surgery on weekdays, and during surgery hours any phone calls for her are transferred to the receptionist who deals with them appropriately. She also has a slot each day when she does house calls. During this period any calls are transferred to the receptionist unless they are from any of her close family (husband, son or daughter) or the receptionist in which case they will be passed directly to her, wherever she may be. Normally she would prefer to receive such calls on her PDA except when she is driving in which case she prefers to receive them on the car phone. Out of hours, when there is no receptionist on duty, any calls from patients should be rerouted to an emergency on-call service. Again calls from family and friends should be put through. If she is at home and the home phone is free, this should be used, if she is in the car, the car phone should be used and in all other cases the call should be directed to her PDA.

One of her maternity patients who is regarded as a having high risk of complications, wears a foetal heart rate monitor (FHR). This sensor detects the heart rate of the unborn baby and can connect to the telecommunications network and send a message when there is a problem. If a problem does occur Carol wants the message to be directed to her immediately wherever she is. The message will contain a string of FHR data. In order to interpret this Carol prefers it to be displayed as a graph – thus the data must be converted to a graph using an appropriate software package, and then, if necessary, to an image that can be displayed on her preferred device.

When Carol is in the office, her preferred device is her desktop. If she is not in her office but her laptop is on, it will direct the message to this. If this is not on, it will direct it to her PDA.

2.3 Requirements

From these three scenarios just outlined the following requirements can be identified.

(1) The user should be able to control to whom a message should be delivered. When a message arrives for the user, the system needs to decide whether or not to deliver it to the user or to deliver it somewhere else. This means that the user needs to be able to specify to whom such a message should be delivered. This will depend on:

(a) The sender of the message. In each of the three scenarios this is the case. In Anna's case, if the message is from the BSL monitor, it must be treated in this particular way. In Bart's case a phone call from his boss has a high priority and is delivered even though he is at home (and probably not accepting any other work calls). In Carol's case phone calls are routed according to whether they are from patients, receptionist, family and friends or from a patient's monitoring device.

(b) The time of day/ day of week. This is relevant to the third scenario. Depending on the time of day/day of week Carol may be doing house calls (weekday) or her own personal activities (weekend).

(c) Context of the user. The most obvious aspect of context is location, and this may clearly be an important factor in deciding whether or not to deliver a message. In the first scenario the locations of potential recipients are used to determine to whom to send the message. In the third scenario the transfers depend on where Carol is – in the surgery or away from it.

In addition to any user-specified preferences, it may also be possible for the system to provide a set of default preferences to assist the user, e.g.
"If the user is on an airplane then …"
"If the user is in a cinema or theatre then …"
"If the user is a scholar and is in the classroom then …"
And so on.

However, location is not the only aspect of context that could be important in deciding whether or not to deliver a message. Another possible attribute is current activity – for example, "If the user is driving a car then …" or "If the user is giving a lecture then …"

Researchers in the area of context management have proposed a variety of other aspects of context, including even emotion (e.g. "If user is feeling depressed and sender is mother-in-law then do not deliver"). However, this will all depend on what context attributes can be provided by the system. In general, as far as this requirement is concerned, any context attribute pertaining to the user that is available in the system should be at his/her disposal in determining whether or not to deliver a message.

(2) If the message is to be delivered to the user, he/she should be able to control which device is used to deliver it. As before the user should be able to specify his/her preferences as to which device should be chosen under what circumstances.

In the first scenario Anna has specified that messages from the BSL monitor should be delivered to her PDA. In the second scenario Bart prefers voice call messages to be connected to his mobile phone unless he is in his car in which case they should be directed to his car phone. In the third scenario Carol prefers the use of her home phone while at home, her car phone while she is in her car and her PDA in all other instances. She also has preferences of PC, laptop and PDA for messages from the patient's monitoring device.

(3) If the message is to be transferred elsewhere, the user should be able to change it. In the case of the message from the BSL monitor to Anna, this must be replaced by a different message before routing it to her parents and/or friends.

(4) If the format of the message is not compatible with the acceptable formats for the device or the preferred format of the user, but it can be converted to an acceptable format, the system should find a conversion process to achieve this. This is particularly relevant in the third scenario where data are being sent by the monitor and need to be displayed in a suitable form. However, it might also apply to the second scenario where the newscast is displayed in one form in the home and another in the car.

(5) If the context of the user or his/her environment changes, resulting in a more appropriate preferred device becoming available, the system should dynamically reconfigure the system to connect the message to the device that is more appropriate/preferred. The obvious context attribute here is the location of the user. In the case of a mobile user, as he/she moves around, his/her environment will change. An example of this from the second scenario is the case where Bart moves from his home to the car while conducting a voice call with his boss. When he gets into the car, the voice call is transferred from his PDA to the car phone automatically.

(6) Finally, a requirement that is implicit in these scenarios is that, where appropriate, the user should have some control over the network that is used. For example, if a land line and a mobile network are both available the user would generally prefer to use the land line since it is cheaper. On the other hand if the user is moving around, the attributes of the service (e.g. Quality of Service) may change. In this case the user may prefer another network that will meet his/her requirements, and one would expect that a pervasive system might dynamically reconfigure to use this.

3. Basic Model

This section presents a model of personalised communications control in a pervasive environment, which is based on the requirements identified in section 2.3.

At the heart of the model is the unit of communication. This may be a message of some form (email, SMS, MMS, etc.), a voice phone call or even a videoconference or broadcast video stream. Any of these will be referred to as a message $M_{i,j}$ where i is the originator of the message and j the intended recipient. The message itself represents the content – a text string, a mixture of text and images, an audio stream or even an audio-video stream. In the case of a broadcast stream, the message becomes simply M_i since it is open to any j to receive it.

From the point of view of the user, the model can be regarded as having three or four main decision or transformation stages. Thus when a new message arrives in the system, it may be subjected to one or more of the following:

(1) Recipient/message transformation. The first decision/transformation stage is aimed at determining whether or not this message is going to the right person and, if not, whether it needs to be changed in any way. In other words this stage performs the transformation R_1 where:

R_1: $M_{i,j}$ -> $N_{i,k}$

where M and N are messages (generally the same), and j and k are recipients.

To illustrate this, consider a few examples. First, suppose that a manager, Bob, is going on holiday but wants to be sure that email from a certain client (Jim) is dealt with quickly in his absence. He might set the system to redirect such email to his secretary (Jane). In other words the system must perform the transformation

R_1: $M_{jim,bob}$ -> $M_{jim,jane}$

Here the message is unchanged but the recipient is replaced. The action is unconditional.

In the case of the third scenario, when a voice call comes in for Carol while she is in her surgery, it is transferred to the reception desk, i.e.

R_1: $M_{i,carol}$ -> $M_{i,reception}$ if loc(Carol) = surgery

Here it is assumed that the call to the reception is essentially the same call as to the doctor in as far as its content is concerned, and hence M remains the same. Thus this transformation again simply replaces the recipient of the voice call.

In the first scenario a message sent from the blood sugar monitor (bsm) to Anna is replaced by a different one to her friends if they are nearby, i.e.

R_1: $M_{bsm,anna}$ -> $N_{bsm,i}$ if unavailable(Anna) & friend(Anna, i) & dist(loc(i), loc(Anna)) < d

(2) Time. Once it has been determined what message should be sent to whom, the next question is when – in other words, should the message be passed to the user now or should it be delayed until later? One way of handling this might be to redirect the message to some form of storage. In the case of a simple text message, it could be stored for later delivery, e.g.

R_2: $M_{i,j}$ -> $M_{i,store}$

or in the case of a voice call it could be redirected to an answering machine, i.e.

R_2: $M_{i,j}$ -> $M_{i,answering\ machine}$

For example, if the user is on an airplane or sitting an exam, and hence he/she is unavailable, these transformations could be used to redirect messages until the user is available once again.

On the other hand, the user might want to delay delivery of a message until a particular time or until a particular event occurs. In the first case the transformation might be:

R_2: $M_{i,j}$ -> $M_{i,j}$ when time = t

while in the second case it is:

R_2: $M_{i,j}$ -> $M_{i,j}$ when event condition

An example of the latter is to use it to set up a reminder message. For example, if the user wants to be reminded to buy milk next time he/she is near the supermarket, the message M would consist of the reminder and the transformation applied might be:

R_2: $M_{i,\ i}$ -> $M_{i,\ i}$ when loc(i) = "supermarket"

This model is not concerned with the actual context management aspects, and assumes that the mapping of location coordinates to meaningful strings can be handled by a context management system. We are also aware of the potential problems that this could raise with privacy but assume that appropriate measures are put in place to protect the user against such problems.

(3) Selection of most appropriate device. Having determined that a message M is to be directed to a user j now, the system must determine what device D_j to use – in other words what device would be the most appropriate for user j. This selection is based on user preferences. If the user has specified a specific device (e.g. his/her PDA) to which messages may be delivered, this device is selected. Alternatively, the user may have specified different devices under different conditions, e.g.

D_j = home phone(j) if loc(j) = "home"
D_j = pda(j) if loc(j) <> "home"

(4) Selection of most appropriate network. In some cases, for example voice calls, there may be different possible network options available to the user regarding the type of network connection used. These may differ in their cost or in the current level of Quality of Service (QoS). For example, the user may have access to GPRS, UMTS, WLAN, … In this case the user may specify the network preferences in a slightly different way, e.g.

N_j = min_cost if loc(i) = "home"
N_j = WLAN if loc(i) = "work"

where min_cost represents the network with the minimum cost from the available networks at any instant.

4. Some Extensions to the Basic Model

This section describes a few extensions to the basic model given in the previous section.

4.1 Conversions

The transformed message $M_{i,j}$ originating from source i needs to be directed to device D_j for recipient j. However, the format of the message may not be compatible with the formats acceptable by the device. If this is the case, one needs to find a conversion process that will convert the message to the format required for device D_j. This may be a single conversion routine or, more generally, a sequence of conversion routines.

Suppose that format($M_{i,j}$) denotes the format of the message and format(D_j) an acceptable format for device D_j. Then this step is seeking a conversion transformation C such that

C: format($M_{i,j}$) -> format(D_j)

4.2 Service Composition

This step essentially links together the message, the conversion transformation process, the device and, if necessary, the network to create a composed service that provides the message on the device for the user. Thus the service composition process is a transformation S such that

S: compose($M_{i,j}$, C, N_j, D_j)

4.3 Service Re-composition

The final step in the process monitors the context of the user while he/she is connected to the message through the device. If the context changes so that a more appropriate device or network becomes available, section 3 step 3 (device selection), step 4 (network selection) and section 4.1 (optional conversion selection) may be repeated as required to recompose the service and provide the message on the new device or via the new network without having to return to the beginning.

Here we do need to distinguish between two types of messages:
(1) Single connection messages (e.g. SMS, email) which, once they are sent by the originator, do not maintain any link with him/her and thus consist of a single composed service at the recipient's end.
(2) Dual connection messages (e.g. voice call) which maintain a link with the originator throughout their existence. Such a message consists of two composed services (one at each end).

For single connection messages only the context of the recipient needs to be monitored and:

S: compose($M_{i,j}$, C, N_j, D'_j) if context of j changes and device D'_j is now best

However, for dual connection messages this process needs to be performed at both ends.

An illustration of this arises in the second scenario. When Bart enters his car, the voice call from his boss is automatically transferred from his PDA to his car phone.

5. Simplifying the Model

One of the first problems that one has to deal with in building any system of this kind is that the set of user preferences can become quite complex. On the one hand one wants to allow user preference formats to be as general as possible to provide as much flexibility as one can to the end-user. This is important if one is going to provide an effective service. On the other hand it is difficult for the end-user to make use of this to specify a complete set of preferences. Indeed, there is an argument that one can never have a complete set of preferences for a particular user as some preferences will inevitably change with time.

To aid the user in this task, the set of user preferences can be broken into subsets in two main ways, namely

(1) Use of roles or identities
It is often possible to identify subsets of preferences based on some way of classifying the user's current situation. For example, if one were to use the notion of roles, the first scenario might be handled as follows. When Carol is in her surgery, one might assume that she is in her "work role" or "professional role". When she is at home she is in her "off-duty role". One might divide these further into "sport role" when she is in the gym or on the golf course, "entertainment role", "holiday role", and so on.

An alternative approach, and that used in the Daidalos pervasive system is to assume that the user has a number of virtual identities, or VIDs [8], and that the user's preferences and current context will determine what VID to use.

This means that instead of having to overload a user preference with additional context information, one can have different preferences stored against different virtual identities (or roles).

For example, suppose that one has the preference rule

when call arrives
if user is at home & (it is a weekend $\vee$ it is a public holiday $\vee$ it is between 6pm and 8am) & caller is a patient
then transfer call to Reception

If the condition

"user is at home & (it is a weekend $\vee$ it is a public holiday $\vee$ it is between 6pm and 8am)"

corresponds to a particular user role ("off duty") or determines a particular virtual identity, one can create a specific instance of the rule, e.g.

when call arrives
if caller is a patient
then transfer call to Reception

and store it with the preferences associated with that user role (“off duty”) or VID.

It is possible that the system may incorporate a learning mechanism that gradually builds up a system of rules or a neural net to determine the appropriate VID or user role to use at any time. In the absence of other alternatives, the user could specify his/her own rules through an appropriate GUI.

(2) Separating user and device/network
Another obvious way in which the user preferences can be partitioned is in terms of the decision as to which user should receive the incoming message (and when) and which device/network should be used for this purpose.

The first decision concerns the user who should receive the message. If the message is directed towards Carol, depending on her current context she may want to transfer it to the receptionist or to the on-call service, or she may want to receive it herself. Likewise, in Anna’s case the message needs to be sent to her parents or friends depending on her and their locations respectively. In the case of a business executive, a message might be transferred to his secretary.

Once this has been determined the second step is to determine which device to use.

6. Conclusion

The aim of this paper was to present a general model developed within the EU research project Daidalos to describe the processes needed to provide users with more control over the delivery of different forms of communications in a pervasive environment. The model has been further explored within the research project Persist (PERsonal Self-Improving SmarT spaces) which is also aimed at the development of pervasive systems.

Three scenarios were presented to illustrate different types of situations and from these a set of requirements was derived. These were used to help define a basic model in section 3. This model was extended in section 4 while section 5 considered how to simplify it.

The model was developed in parallel with the implementation of two slightly different pervasive systems prototypes in which some aspects of communications control were implemented and demonstrated. The first scenario from section 2 was demonstrated in the first prototype using a simple form of SMS messaging (using SOAP). The second scenario was part of a major demonstration and was based on VoIP (Voice over IP) services (using SIP). On the other hand format conversions were not included and recipient substitution was limited. Both composition and re-composition have been demonstrated. Network preferences were also implemented although this was handled at a low level in the system as opposed to the high-level handling of composition. The ideas are being studied further in another EU research project, Persist, which aims to develop a more powerful model for pervasive systems that can support a more truly ubiquitous system.

Acknowledgements

The research leading to these results has received funding from the European Community through the research project Daidalos (Sixth Framework Programme) and the ideas have been further studied in the project Persist (no. 215098) in the Seventh Framework Programme. The authors also wish to thank all colleagues in the Daidalos project developing the pervasive system. However, it should be noted that this paper expresses the authors’ personal views, which are not necessarily those of the Daidalos consortium. Apart from funding these two projects, the European Commission has no responsibility for the content of this paper.

References

[1] T. Kindberg, & A. Fox, System Software for Ubiquitous Computing, *IEEE Pervasive Computing*, *1*(1), 2002, 70-81.
[2] N. Kohtake, K. Matsumiya, et al., Smart Device Collaboration for Ubiquitous Computing Environment, *Proc. Ubicomp 2003 Workshop: Multi-Device Interfaces for Ubiquitous Peripheral Interaction*, Seattle, Washington, 2003.
[3] M. Satyanarayanan, Pervasive computing: vision and challenges, *IEEE PCM*, *8*(4), 2001, 10-17.
[4] G. Abowd, & E. Mynatt, Charting past, present and future research in ubiquitous computing, *ACM Trans on Computer-Human Interaction*, *7*(1), 2000, 29-58.
[5] S. Abu-Hakima, R. Liscano, & R. Impey, A common multi-agent testbed for diverse seamless personal information networking applications, *IEEE Communications Magazine*, 36(7), 1998, 68-74.
[6] M. Weiser, Some computer science problems in ubiquitous computing, *Comm ACM*, *36*(7), 1993, 78-84.
[7] M. H. Williams, N. K. Taylor, I. Roussaki, P. Robertson, B. Farshchian, & K. Doolin, Developing a Pervasive System for a Mobile Environment, *Proc. eChallenges 2006 – Exploiting the Knowledge Economy*, Barcelona, Spain, 2006, pp. 1695 – 1702.
[8] J. Girao, A. Sarma, & R. Aguiar, Virtual identities - a cross layer approach to identity and identity management, *Proc. 17th Wireless World Research Forum*, Heidelberg, Germany, November 2006.