A Personal Distributed Environment for Future Mobile Systems

J Dunlop, R C Atkinson, J Irvine & D Pearce

Mobile Communications Group
Dept. Electronic & Electrical Eng.
University of Strathclyde
Glasgow, UK, G1 1XW
Email: j.dunlop@eee.strath.ac.uk

ABSTRACT

A Personal Distributed Environment (PDE) embraces a user-centric view of communications that take place against a backdrop of multiple user devices, each with its distinct capabilities, in physically separate locations. This paper provides an overview of a Personal Distributed Environment and some of the research issues related to the implementation of the PDE concept that are being considered in the current Mobile VCE work programme.

I. INTRODUCTION

The Personal Distributed Environment (PDE) encompasses a User perspective of multiple devices (both local and remote) accessing multiple services via multiple networks, all of which can be changing dynamically. It encapsulates the concept that coverage is not necessarily universal but may occur in islands which may or may not be inter-connected. This implies that a particular session may not be continuous but is commenced or continued whenever the user is within range of service delivery mechanisms which may include broadcast delivery, mobile cellular networks, low power personal ad-hoc radio networks, wireline networks, etc, as shown in Figure 1.

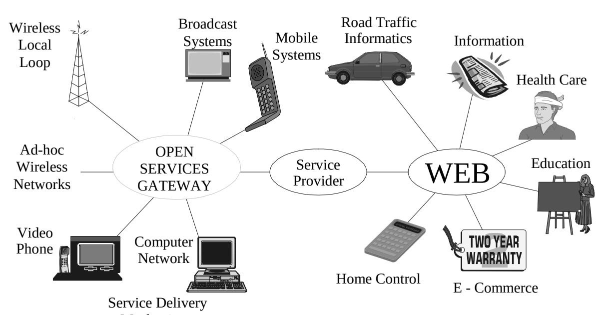

Figure 1: PDE Schematic

The main issue to be considered in relation to the PDE concept is the provision of service discovery, negotiation and delivery options for multiple service providers within an that may contain one or more "Distributed Terminals", available to each user. The ad-hoc nature of the PDE as location changes is a major issue in defining the capabilities and functionality of the distributed terminals.

The Personal Distributed Environment concept is key in the creative interworking of wireless and broadcast technologies for enhanced service provision. The essential issue is that the user terminal is no longer a single device (which includes multi-standard devices) but rather a "distributed" set of devices which have a range of capabilities. Of particular interest is the consensus that interworking of wireless and broadcast technologies for service provision will be accompanied by a user-centric emphasis. The user will be empowered to negotiate his/her service facilities and tariffs at the point of delivery and will be able to control access and quality for distributed terminals whose characteristics will be location and time dependent. These capabilities must, however, be user-friendly.

The PDE concept allows a number of physically separated devices to access the content stored on each through a pervasive network. The PDE may therefore be considered as a combination of local devices, connected using PAN technologies such as Bluetooth and ZigBee, and remote devices using a variety of heterogeneous networking options such as fixed line, cellular and WLAN technologies. Some of these devices may join and leave the PDE in a dynamic way and issues raised by this dynamic reconfiguration are addressed later in this paper. In principle the PDE user is able to use any device within the PDE to access content stored on any other, subject to security constraints, device capabilities, and connectivity issues. One such constraint which needs to be accommodated, for example, is that devices connected through wireless links may have only intermittent connectivity. In fact intermittent connectivity may be extended to entire sub networks, such as Personal Area Networks (PANs), which may be connected using short-range radio links and this extends the dynamic nature of the topology.

Since each device may utilise different access technologies (e.g. fixed-line, cellular, WLAN, or combinations thereof), and each device may have different capabilities (e.g. processing power, battery life, video capabilities, local storage), the PDE concept seeks to harness these capabilities in novel ways to permit a new and diverse range of services to be delivered to the user. The PDE concept primarily puts the user in control of his/her communications environment.

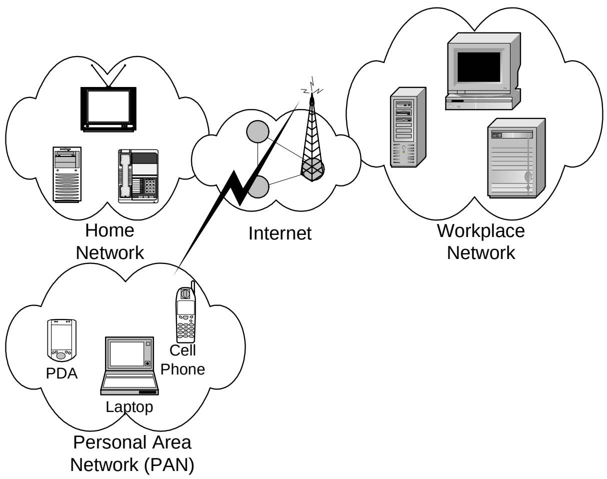

Figure 2: PDE Sub-networks

The PDE is concerned with communication between its constituent devices and also with how a device, or combinations of devices, communicate with various external players (network operators and content providers). In line with recent research (DRiVE Project [1], COMCAR Project [2]) aimed at utilising distinct networks (DxB, cellular, WLAN, PAN, etc.) to deliver a new range of services to users, PDE permits devices with different access technologies to receive bit streams from a number of networks and combine them to form advanced services. For example, a user may receive a video stream over DVB-T and subtitles over cellular. Essentially, the PDE enables devices to operate together in a co-ordinated fashion.

The paper is arranged as follows. Section II discusses the requirements to facilitate internal communication within the PDE. Section III discusses the issues of device location against a backdrop of multiple devices with a range of capabilities. Section IV describes the role and necessity of feature discovery in the overall PDE architecture with particular emphasis focused on the different approaches proposed to resolve this problem. Section V discusses security considerations within the PDE, and Section VI presents some suggestions on PDE development.

II. INTERNAL CONNECTIVITY

A typical selection of sub networks which could form the elements of a PDE are shown in figure 2. In this figure the PAN, for example, may contain a cellular phone (e.g. UMTS), a WLAN-enabled laptop and a PDA. As these units are part of a PAN it is likely that all three devices are Bluetooth enabled such that any one device can communicate with each of the others. With this capability a data call, for example, could be set up to the cellular phone over a UMTS network or over a WLAN network to the laptop with the last hop over Bluetooth to the handset. In this case two routes would exist to the handset. The situation would be complicated further if the laptop were cellular-enabled also in which case there would be two routes to the handset via cellular networks e.g. directly over a cellular link to the handset, or indirectly over a cellular link via the laptop and a Bluetooth link.

In the event that two or more devices in the same PDE sub-network are able to use the same technology (i.e. UMTS, GPRS, WLAN, etc.) to access the wider network, then some coordination is necessary within the PDE to avoid several devices continuously monitoring the same air-interface. Such a coordination is also required to determine the most appropriate route to a specific device within the PDE.

III. DEVICE LOCATION ASPECTS

Services may be provided to entities within the PDE by devices located within the PDE (intra-PDE services) or from devices or service providers located outside the PDE (extra-PDE services). In order to support intra-PDE services, the architecture must provide mechanisms for topology management which involve both local and remote devices and must therefore be able to accommodate dynamic reconfiguration of the PDE. For example devices (or sub-networks) may become unreachable from other elements of the PDE through one access technology yet still reachable through another. Alternatively, devices and sub-networks may be reachable via a number of different access technologies, each with different tariffs, billing mechanisms and QoS; it may be advantageous to select one technology in preference to another depending upon particular circumstances.

Thus topology management of the PDE must consider *application layer connectivity* of devices within the PDE over a range of radio and fixed access technologies. There is a requirement for functionality to store partial routes between the constituent devices of a PDE.

Extra-PDE services include those where an application or session originates or terminates at a device outside the boundary of the PDE. Extra-PDE services place additional requirements on the PDE architecture in terms of feature/service discovery (discussed in Section IV) and device location.

Whenever an outside agency wishes to establish a session with a user who is communicating from within a PDE, a first point of contact is required. This situation arises because the PDE may include several devices with functionality that can support the session, which may not be local to the user. One possible solution to this problem is the establishment of a Device Management Entity (DME) that can facilitate the session set-up by determining the appropriate device(s) within the PDE to handle the session.

Device location knowledge within a PDE is required so that an incoming service may be directed to a device capable of accepting that service and which is accessible to the user. This also allows identification of several service termination points (forking) .

Approaches to device location indication include proxy-based and redirect-based; both of which can be implemented using The Session Initiation Protocol (SIP) [3], for example. With the former, the PDE contains a proxy server that, on receipt of session establishment requests, determines the most appropriate device (based on device capabilities, preferences and knowledge about device and user location) and forwards session establishment requests to that device. The device in turn signals its responses back to the proxy; the responses are subsequently relayed to the source. With the latter approach, session establishment requests are received at the redirect server. The redirect server identifies the most appropriate device to support the service and returns this information to the source. The source then sends a session establishment request to the target device. With the redirect-based approach, the server acts in many ways like a Domain Name Server (DNS); the redirect server merely returns a handle to the most appropriate device. Thereafter session set up is handled on an end-to-end basis.

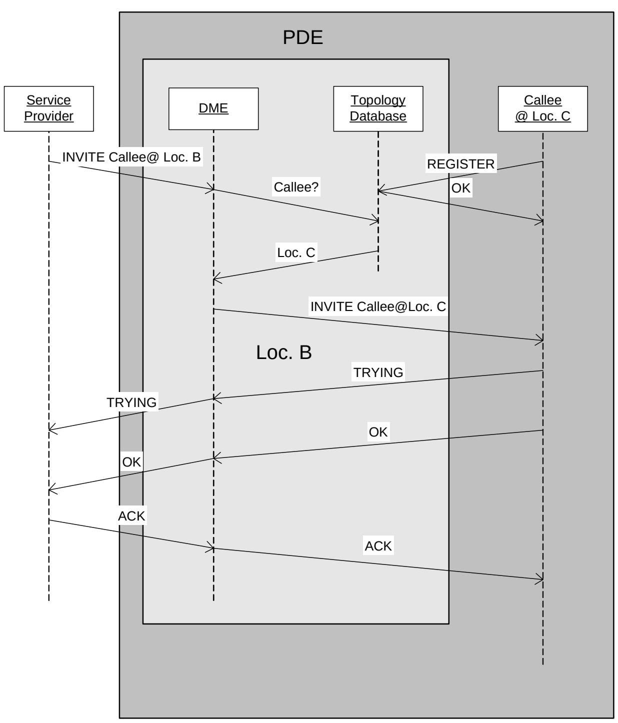

Figure 3 SIP Proxy Server Signal Flow

Signalling between a service provider and a PDE is illustrated in Figure 3 where proxy mode is employed. The end party (callee) registers his/her location (location C) via a location update to the topology database. A third party service provider attempts to communicate with the user at the location of his/her DME (at location B). The DME interrogates the Topology Database for the physical location of the user. The DME then forwards request to the user on behalf of the service provider.

In order to facilitate both intra-PDE and extra-PDE communication, a topology database is required. The topology database keeps track of reachability information to the various PDE devices across the various access technologies. The database can be queried to obtain the location of particular devices; in this context location is not a geographical attribute rather it is information on which particular network a device is located. The topology database functions in some respects in a similar way to the Location Server required by the SIP [3].

It is apparent, therefore, that the devices that comprise the PDE must register their location with the topology database from time to time. In the case of a cellular phone location information may be simply identity of the network to which it is currently attached (this takes into account roaming agreements), in the case of a WLAN-enabled laptop this information would be the particular IP sub-network to which it is attached, and in the case of a home PC or a workplace server the information would be a static IP address.

Given that the topology database is dynamic and reflects the dynamic nature of the PDE topology, it is necessary to take steps to minimise the signalling overhead associated with frequent location updates. Possible mechanisms to address this include asynchronous registration and establishment of a distributed database. Asynchronous registration refers to the ability of different devices, depending on their access technology, to update their location record on different time scales. For example, a WLAN-enabled laptop may require to update its location record whenever it is activated and detects a radio port; alternatively a PC with a fixed location and a static IP address may never require to update its record.

As indicated, the PDE can be decomposed into sub-networks and within each sub-network a nominated device may act as a gateway to all other devices for each access technology. Therefore, in some cases it may be necessary only for the topology database to record that a particular device, for example the Bluetooth-enabled PDA in figure 2, is accessible through the cellular phone. As the PDA moves around its environment it need not update its location record as long as it is attached to the cellular phone. Within a sub-network, each of the nominated gateway devices will require reachability information for the other devices within that sub-network, therefore only partial reachability information for some devices in the

PDE may be stored in the topology database; the remainder can be stored in local gateways. Thus the topology database can be regarded as a distributed database and only local changes need to be reported locally.

IV. FEATURE DISCOVERY ASPECTS

Feature discovery is an essential element in the operation of the PDE. It enables delivery of content-rich multimedia services in a format appropriate to a specific device and the network to which it is attached. Furthermore, feature discovery enables opportunistic communication; devices within the PDE with ad hoc networking capabilities (such as those contained in a PAN) would be able to communicate effectively with other ad hoc networks or enabled devices that they encounter. For example, a Bluetooth-enabled laptop would be able to make use of a large display that is also Bluetooth-enabled.

Feature discovery is a mechanism whereby foreign devices and services can determine the capabilities of a particular device through interrogation. For example, a device with a liquid crystal display (LCD) should be able to provide information on the resolution and colour range of that display upon request. Service discovery is an analogous procedure aimed at identifying the services operating on devices. Together service and feature discovery make it possible for multimedia services to be tailored to suit the end device.

There are several approaches to service delivery based on the capabilities of devices contained within a PDE. In the ad hoc networking community, where service discovery has evolved, two main approaches have been developed: end-end discovery and directory-based discovery. With end-end discovery, any two devices that wish to communicate exchange signalling information to describe their capabilities; an example of this approach is the Service Discovery Protocol (SDP) [4] employed by Bluetooth. With the directory-based approach, devices register their capabilities with a directory service; any device that wishes to communicate with a target device interrogates the directory service (on another device) for the capabilities of the target device. An example of this approach has been adopted by the Jini protocol [5].

In the telecommunications community, as part of the H.323 suite of protocols, H.245 [6] describes an end-end method of specifying the receive and transmit capabilities of a device based on a capability table held within the device. For example, a device may list that it can handle audio in G.711, G.723.1, and G.728 formats. This information is exchanged after a connection between two devices has been established. In the Internet community, under the auspices of the IETF, the SIP protocol takes a prescriptive approach. With the Session Description Protocol [7] associated with SIP no feature discovery is conducted; rather as part of the session set-up procedure the capabilities of a device required to receive a service are listed in order of preference, e.g. G.711 then G.723.1, etc.
Since different sub-networks of the PDE may have different capabilities (power and bandwidth), different techniques may be more applicable to different sub-networks. Thus, not only must the PDE architecture be able to support a range of discovery protocols, but it must be able to support interworking of distinct discovery protocols. Interworking of discovery protocols has been a topic of research within the Open Services Gateway Initiative (OSGi) [8].

V. SECURITY CONSIDERATIONS

The power and flexibility of the PDE introduces a number of significant security concerns. These include the fact that the identity of the PDE must be verified, accesses to information and services must be authorised, and finally any information transfers within the PDE must be secure.

Identity is an important consideration in the PDE context. Since the PDE is distributed over a number of devices, traditional solutions such as SIM cards, using secure hardware to match an identity with a single device, are no longer appropriate. Also, the PDE topology is dynamic, and can be instantiated in several physically separate locations on devices which may not belong to the user, so a software solution is required. However, it is important part of the PDE is not able to act autonomously relative to the user in an unauthorised manner, and that communicating entities and service providers can trust the PDE, so some sort of centralised Security Information Register will be required.

Authorisation of devices within the PDE will need to be managed, with security policies defined for different devices, for events such as transferring information or running code on different devices, as defined by integrity, privacy or commercial requirements.

PDE security services will also deal with content security, both of user data and data from service providers. To meet the trust requirements of the different access technologies with which the PDE may interwork, the security service is likely to have a number of operator domains, handling security with specific networks.

VI. PDE DEVELOPMENT

It is clear that research is required into the structure and composition of the topology database; signalling overhead being a prime area of attention. The influence of devolving topology information to local sub-networks on signalling overhead is a significant issue.

Where the sub-network contains devices that can connect to the same access network, mechanisms must be developed to determine the most appropriate device to act as a nominated gateway.

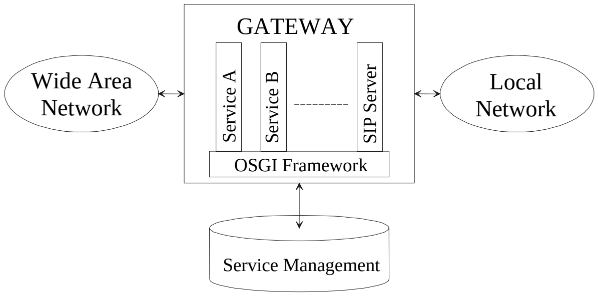

Figure 4: OSGi Gateway

Several aspects of the PDE and candidate technologies have been outlined to facilitate service signalling. Inherent within these solutions are several assumptions which may not hold true in a PDE. The SIP solution for example assumes all devices and services are represented in a familiar manner, allowing for services to be compared and thus properly selected and addressed from a SIP database. The SIP and H.323 suite of protocols were developed based on characteristics of devices found in the telecommunications and Internet worlds. Many PDE devices are not designed as such and are therefore incapable of interacting directly with a SIP server, and they may not able to write to or read from a H.245 table. Additionally, the various discovery protocols, a few of which were briefly mentioned in section IV, use different data representation formats such as XML and ASN.1, thereby complicating service selection.

The Open Service Gateway Initiative (OSGI) outlines certain principles which may be instructive in overcoming this problem. OSGI creates open specifications for the managed delivery of multiple services over wide area networks to local networks and devices. The OSGI forum has defined a scheme where all devices regardless of network technology and functionality can be represented in a transparent manner at various levels of abstraction. The Gateway consists of a framework, upon which service bundles are added to create functionality within the gateway, as and when they are needed. This can provide a significant element of the DME at a local network level. OSGI allows for the development of a service which will represent the devices (to the outside world) in such a scenario. In addition an option exists for the SIP server to be implemented as a service bundle on the gateway, as shown in figure 4, or to be managed remotely.

This is convenient as:

- The gateway can recognise various devices and represent them in a common way
- The gateway can register these devices with the SIP server, which can then handle service requests

The OSGi gateway was initially created to allow for various devices in separate home networks, using different discovery protocols to be able to interact transparently with each other and the outside world. Central to the implementation was a gateway device through which all devices communicated with the outside world. While the implementation is not consistent with that of PDE the principles are quite relevant. In this respect it is worthwhile to note that even within the OSGi community, there have been projects aimed at porting OSGi to mobile devices and vehicles.

The potential for incorporating the OSGi gateway into the overall PDE architecture is an area of investigation within the MVCE research programme which will also assess the potential of SIP and other protocols for signalling device location information to interested parties

VII. ACKNOWLEDGEMENTS

The work reported in this paper forms part of the Core 3 Research Programme of the Virtual Centre of Excellence in Mobile & Personal Communications, Mobile VCE, www.mobilevce.com whose funding support is gratefully acknowledged. More detailed technical reports on this research are available to Industrial Members of Mobile VCE.

REFERENCES

1. L Christodoulides et al., “DRiVE towards systems beyond 3G”, in Proceedings of SCI 2001; Orlando; 22.-25. July 2001.
2. COMCAR: Communication and Mobility by Cellular Advanced Radio, http://www.comcar.de/overview.pdf
3. M Handley et al., “SIP: Session Initiation Protocol”, IETF RFC 2543
4. Bluetooth SIG, Specification of the Bluetooth System V 1.0b, Part E, December 1999, pp 323 – 384.
5. Sun Microsystems Inc., Jini: An Architectural Overview, http://www.sun.com/software/jini/whitepapers/architecture.pdf
6. Control Protocol for Multimedia Communication, ITU Recommendation H.245, ITU.
7. M Handley et al., “SDP: Session Description Protocol”, IETF RFC 2327
8. P Dobrev et al., “Device and Service Discovery in Home Networks with OSGi”, IEEE Commun. Mag., Vol. 40, No. 8, Aug. 2002, pp 86-92.