Realms as Entity Location Services

A study of object orientation for the conceptual structure of the Internet

Oliver F. J. Snowden

for

The Nuffield Foundation

Summary

New methods for organising the Internet need to be investigated and a new interface concept called **Realms** explored. It is hoped that Realms will enable access to a variety of devices not typically accessible over the Internet.

This is needed since the Internet is rapidly expanding and huge numbers of devices are expected to be connected simultaneously (e.g. electricity meters, TVs, washing machines, individual lights, pipeline sensors). This introduces many difficulties, especially in referencing and organising such vast numbers of devices. Without efficient organisation current priority services, such as telephony, will be degraded.

We have researched the following topics:
- Current and potential Internet devices
- Current and planned structure of the Internet
- The Internet as an Object Oriented distributed database
- The Realm concept with respect to the conceptual structure of the Internet.

In addition we have
- Ascertained any potential increase in productivity from Internet users using this concept
- Conducted interviews with potential end-users and experts
- Studied realm interfaces,

and written critical analyses of
- The Realm concept for the Internet structure
- The Realm concept for widening access.

The results of this study suggest that the scope of the devices connected to the Internet is likely to increase and that there are mechanisms which can facilitate the increased number of devices, both physically and virtually.

Object Orientation (OO) was explained briefly to help highlight the benefits it can bring to users and computational systems. This provided a grounding for the Realms concept: an attempt to apply the OO technique to Internet systems. A number of potential human and computation benefits were explored which were primarily focused around human context and computation standardisation.

These ideas were proposed to both potential end users and University of Southampton networking experts. The response from the limited number of end users was largely positive, although it was evident that their desire for Realms may have been influenced by their desire for an unfulfilled need, as opposed to the Realm structure to bind the services. However, the networking experts questioned largely gave a negative response to the Realms concept. They felt that the concept was primarily too complex and stated that there were other approaches to solving the problems associated with the benefits that Realms could provide.

From a general Realms foundation and the response from the interviewees, the research focused on realm interface. The Dewey Decimal System was briefly studied and it was found, due to copyright and the associated financial obligations, it was unlikely this system could be used for realm

categorisation. Other categorisation facilities were mentioned, although this subject was too extensive for a research project of this size. Further to categorisation, the realm interface was discussed in terms of current Internet systems (*such as DNS*). JXTA was mentioned as a system capable of supporting such a realm system.

The remaining study focused on criticising the various aspects that the Realms concept covered. It was evident that there were significantly adverse human factors associated with content categorisation and that there were limitations in the management of peer-to-peer systems. Universal Description, Discovery and Integration (UDDI) was briefly studied as an alternative approach to a peer-to-peer architecture.

Widening access was the final aspect to be studied. The critique factors were taken from "sensor web" documentation and it was largely concluded that if the Realms concept, or any similar concept, were to be widely used it would be likely that there would be increased access to information.

Acknowledgements
I would like to thank my research supervisor, Quintin Gee, for not only assisting with this research, but also for the encouragement throughout my studies. I would also like to show my gratitude by thanking everyone who participated in the research, especially those who participated in the interviews. Finally I thank the Nuffield Foundation for funding this research under their Science Award Scheme for Undergraduate Research Bursaries, 2004.

Contents

Introduction........ 5

Icons Used in This Document........ 6

Section 1: Analysis of current and potential Internet devices........ 7

Section 2: Analysis of current and planned structure of the Internet........20

Section 3: Investigation of the Internet as an Object Oriented distributed database........ 29

Section 4: Investigation of the Realm concept with respect to the conceptual structure of the Internet........ 38

Section 5: Interviews with potential end-users and experts........48

Section 6: Study of realm interfaces........59

Section 7: Critical Analysis of the Realm concept for the Internet structure........ 69

Section 8: Critical Analysis of the Realm concept for widening access........ 81

Project Summary and Conclusion........ 89

Introduction

The number of Internet addresses will rapidly increase over the next few years as the new Internet Protocol (IP) standard is introduced (Ipv6).

Abilene (2003) states that IPv6 will vastly increase the number of Internet addresses from a mere ~ 4.3×10^9 (~*4 billion*) to ~3.4×10^{38} (~*340 undecillion*), thus enabling more devices to connect to the Internet simultaneously. *For example, washing machines could be connected to enable the manufacturer to diagnose problems remotely.*

The domain of interest is how a specific device can be located, and how *like* devices should be conceptually organised and accessed. The project will focus around Object Orientation and peer-to-peer collaboration rather than more traditional approaches. The rationale for the community focus is that it could help businesses create cheaper computational systems, promote the establishment of virtual marketplaces, and provide human understandable context to computation systems.

A hardware device refers to a physical device such as a telephone, whereas a virtual device refers to a system implemented in software that performs a virtual function. This could range from a virtual film to a virtual medical doctor.

Icons Used in This Document

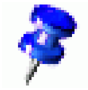

The “Note” icon marks extra information that the reader may find useful.

The “side thought” icon indicates information that provokes further thought, which is typically outside the scope of this document.

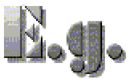

The “example” icon is identifies significant examples.

Pin and light bulb image source: OpenOffice.org

Section 1: Analysis of current and potential Internet devices

Internet Device aliases: Internet appliance, net appliance, smart appliance, information appliance.

It is important to draw a distinction between an Internet device and an Internet appliance. Whatis.com (2004) defines an IA as "*a machine designed for a specific function that also has a built-in Web-enabled computer*" and gives the following examples: E-mail devices, personal digital assistants (PDAs), web-enabled refrigerators. It is thus a device with a specific application. However, Internet Device is used in the document as a more general term to include any device that accesses the Internet/Intranet.

1.1 Computing Development

Weiser (1996a, p.2), perceived by many as the father of ubiquitous computing, identified three "waves" of computing:

- *The first wave of computing, from 1940 to about 1980, was dominated by many people serving one computer.* – Mainframes
- *The second wave, still peaking, has one person and one computer in uneasy symbiosis, staring at each other across the desktop without really inhabiting each other's worlds.* – PCs
- *The third wave, just beginning, has many computers serving each person everywhere in the world. I [Weiser]call this last wave "ubiquitous computing" or "ubicomp".*

Weiser (1996b) described ubiquitous computing as the age of calm technology. He stated that virtual reality puts people inside a computer-generated world whereas ubiquitous computing forces the computer to live out here in the world with people, and is therefore a very difficult integration of human factors, computer science, engineering, and social sciences.

1.2 Current Devices

Dingli (2001) provides four broad categories of consumer Internet device. We have identified two additional categories, "Other (Specialist)" and "Virtual" Devices as being useful to our approach.

1.2.1 Personal Computer (PC) – Fat Internet Client

A device that is typically found in homes and offices that is used to provide a user with a vast range of facilities. There are two generic types of PC, desktops and laptops, with the latter designed specifically to overcome the mobility constraints of traditional desktops.

PCs can accommodate a range of input/output devices (*e.g. keyboards and visual display units*), provide ample processing and storage capacity for personal use, and are one of the most cost effective forms of computing.

1.2.2 Web Enabled TV – Thin Internet Client

Dingli (2001) states that this method of accessing the Internet has been popular in America for several years. NTL (2004) show that TV Internet access is now provided by the major UK cable and satellite companies, NTL and Sky.

Web-pages are made available by using standard modems and communication networks, such as POTS (Plain Old Telephone System) and Cable Television networks for Sky and NTL respectively. The web-pages can be displayed on traditional TVs. The quality (*screen resolution*) of the standard TV, however, is limited. This causes a number of problems including that of maintaining legibility. Traditional web-page buttons and text below 18 point (*size*) can become "fuzzy" on a standard TV screen. NTL (2002) provide guidelines to help overcome the problems, although it is evident that this is a significant limitation for the use of the web on TVs today. Despite this limitation, there are clear uses, especially when integration with interactive content is important (*e.g. interaction with a live TV game show*).

1.2.3 Personal Digital Assistants (PDAs)

Aliases: Palmtops, hand-held PCs
Bocij *et al.* (2003) states that these devices are typically no larger than a pocket calculator and have only a limited functionality (*e.g. address book, appointment scheduler and note taking*). However, many provide ample processing power to provide a range of services such as word processing and Internet access.

For example, PalmOne Inc. (2004) provide a web-browser and e-mail client for use when the device is connected to the Internet. Input is accomplished through a number of means; touch screens, miniaturised keyboards and microphones are common. However, after speaking to an experienced end user, we identified a number of problems. (Snowden (2004))

- Hardware is far slower than PCs, the screen is small and makes some aspects difficult to navigate (*such as standard web pages*). Further to this, memory is very expensive compared to PC counterparts.
- Software is typically limited in features, is often supplied by one provider and is expensive (*proprietary software*). This means it can be costly to "buy" functionality. An exception to this is the Linux-based Sharp Zaurus, which is able to use popular Open Source software.
- Internet access can be difficult to achieve. Specialist cables and phones are typically required to establish an Internet connection; although 802.11 and 802.15 wireless standards have been implemented in some of the newer PDAs to make the devices more versatile.

1.2.4 Mobile Phones

Alias: Wireless Internet Device.
These devices range in size, although 3G phones are currently comparable to PDAs. Mobile phones are able to provide many of the traditional PDA features (*e.g. Scheduler, Calculator etc.*) and have access to communication networks with good geographical coverage.

1.2.5 Other Devices (Specialist)

There is a growing range of Internet Devices which are available today that don't fit the categories listed above. These devices range in popularity, but haven't been as popular as the devices listed above.

- ***Portable MP3 players***

  Devices, such as the Apple Ipod, see Apple (2004), are becoming increasing popular. A user is able to use relevant software to download music (*legally*) from the Internet directly onto the MP3 player. According to the BBC (2004), Apple Inc's iTunes service has been incredibly popular, with approximately 85 million song downloads within the first 15 months of opening. The service has been available in Europe since June 2004.

- ***IP Phones***

  These devices enable voice and, optionally, video communications between two or more parties. Their use has been limited due to the Internet quality of service (QoS) constraints. Tanenbaum (2003, p.398) discusses various techniques for achieving good QoS, although he notes that good QoS can only be achieved across the Internet if these techniques are employed. After discussion with a Cisco network engineer, Howarth (2004), the QoS techniques that Tanenbaum has provided have not been fully implemented. Further to this, Howarth states that Internet usage is incredibly volatile, citing that any sudden influx of use has a significant impact on QoS. He states that the Internet as a whole is not ready for such activity and has left Voice over IP largely to the confines of businesses that have dedicated lines to help maintain QoS.

- ***Digital Picture Frame***

  This is a specialist device that allows the user to display a "slide show" of pictures and, in some cases, a video in a traditional picture frame. This can be used in homes and offices, for business or pleasure, and can be set up to display any picture (*e.g. pictures of stock market results, advertisements, wildlife images*).

  These devices can be made available on-line, allowing remote administration of the pictures. The images can be controlled by authorised personnel or, optionally, the public. An example of this type of product is provided by Digi-Frame Inc. (2004).

- ***Network Webcam***

  This is a video camera that provides a streaming video feed accessible via the Internet or an intranet. This type of device is popular with a vast range of users, e.g. radio hosts (*for greater public interaction*) and hoteliers (*to display tranquil surroundings*).

1.2.6 Virtual Devices

Alias:soft-devices.
The above documentation listed physical devices, yet there is strong evidence that many of the devices are converging to offer standardised "services". It is these services that one can name "virtual devices" – systems that provide an interface to a particular service accessible via the Internet.

Current examples of virtual devices:

- Childlocate (2004) – this enables a parent/guardian to locate their child/children using mobile phone technology via the Internet.
- Napster (2004) – similar to iTunes, this provides the user with an interface to millions of soundtracks.
- Google (2004) – this enables a user to search many different aspects of the Internet, predominantly web-pages.
- ScienceDirect (2004) – this enables a user to search through and gain access to over 1500 journals.
- Xdrive (2004) – this enables a user to have an on-line filestore (*hard-disk*), enabling a user to access saved data anywhere in the world with Internet access.

1.3 Scope of current Internet Devices/Product Scope

It is evident that current devices are, by and large, still part of Weiser's "second wave" of computer systems, where the user commands the computer actions via clumsy desktops. They are predominantly pervasive (*device focused and user led*) with few systems/devices that actively serve a human.

Dertouzos (2001) states that "*pervasive computing is easy. It's what we already have, only more of it...The Information Revolution won't fulfil its promise until we stop thinking as though we're still in the Industrial Revolution*". This view suggests that most of the current devices are simply tools that are designed to make what were manual tasks, easier, faster and more reliable. Dertouzos insinuates that mankind cannot progress by merely applying pervasive computing to more and more situations, which is what the current trends show.

The perceived lack of progression in this area can be partly explained by the relatively recent mass use of the Internet, and the corresponding recent investment by companies into Internet technology providing the information ether. However, there have been significant technological advancements which are laying the foundations that Weiser, in IEEE Computer Society (1998), once predicted. One of the most significant aspects was the advent of wireless communications standards such as

802.11 (Wireless LAN) – enables TCP/IP to be utilised by wireless devices.

802.15 (Bluetooth) – enables small devices to be interconnected, e.g. Mobile phones.

802.16 (Wireless Broadband) – this technology enables a wireless local loop.

These open standards are allowing all electronic manufacturers to make "ubicomp" products that are both interchangeable and interoperable.

For example: It is currently possible to create a wireless LAN consisting of 802.11 devices from different manufacturers that work seamlessly together. This applies in a similar fashion to Bluetooth devices and Personal Area Networks (PANs).

It is important to acknowledge the importance of ALL open standards, and not just those listed above. Standards such as TCP/IP and HTTP all play their part in the global network.

Home/Office wireless devices are highly likely to be ubiquitous in the not so distant future. It may be possible to extend the Internet architecture to exploit these wireless LANs.

> Example: Mr Smith has a wireless access point (AP) that he allows to communicate with a nearby neighbour's AP (and *vice versa*). In turn, the neighbour permits access to another neighbour. This extends the network architecture at relatively little cost (and relieves the dependence on some proprietary networks).

Request for Comment (RFC) no. 1876 provides a means for expressing location information in the Domain Name System (DNS). This could be used to enable routing based on location information.

1.4 Future Direction

Different organisations have taken different approaches to ubiquitous computing. It has become apparent that there are two strands, shown in Figure 1.1.

	Pervasive Computing	***Human-Centric Computing***
Objective	To create a device that is constantly connected, portable, and available.	To create a solution so that the human is always connected, portable, and available.
Focus	Devices	Humans
Method of interaction	Work with the device for the information	Speak to the solution for the information
Current Projects	IBM's Project Blue Carnegie Mellon's Project Aura	MIT's Project Oxygen

Figure 1.1

Source: Schwarz (2001, p.1)

Schwarz (*loc.cit.*) suggests that pervasive computing technology is likely to continue to be the focus of the consumer and organisations, since human-centric computers require a shift in the way computers are conceptualised.

This can be explained by the relatively recent adoption of globally agreed open standards, that interconnected devices rely upon. One can deduce that Human-Centric Computing relies upon the interconnection of devices, but communication standards (*and corresponding technologies*) are still under heavy development (*e.g. IEEE: 802.11i, IETF: Mobile IPv6 etc.*). Consequently, it is only logical for developers to primarily focus on devices and their interconnection to establish a foundation for Human-Centric Computing.

1.5 Examples of Future Devices

1.5.1 Home Appliances

There is mounting evidence that the Internet will be used to support some traditional home appliances such as fridges and washing machines (*Pervasive Computing*).

Electrolux (2004) provides details of its prototype "Screenfridge" (*see Figure 1.2*) that offers the following facilities:

- Internal Communication – *allows a user to leave internal messages.*
- External Communication – *provides e-mail, web and potentially telephone and video conference facilities. Food items could be ordered, and bills paid.*
- Food Management – *the device could automatically download details on how to prepare food stored in the fridge*.
- News, Radio and Home Security – *The Screenfridge is equipped with a TV, radio receiver and surveillance system access points*.
- Digital Cook book – *The Screenfridge also features a powerful cookbook with hundreds of recipes which could potentially be connected to an Internet database* (*enabling updates/amendments etc.*)

Electrolux Screenfridge

Figure 1.2
Source: Electrolux (2004)

1.5.2 Radio Frequency Identification (RFID)

A significant amount of research suggests that RFID is an extremely promising technology. It is concerned with small devices that are able to broadcast wireless signals. According to Hornby (1999, p.3), the devices consist of a silicon chip that is "bonded" onto an aluminium antenna. They are cheap and small enough to be disposable, and can be attached to inanimate objects for identification purposes.

Foroohar (2004) states that some retail stores are beginning to use them as advanced bar codes on individual products (*within the store*) and batches of products (*during product transit*). This technology essentially reduces the staffing overhead required to identify items, because a scanner isn't required to be physically placed next to each item. The wireless transmissions enable "hidden"items to be identified (*e.g. items within a box*). If adopted, this technology could enable shoppers to place items into a shopping bag, leave the store and provoke automated systems to take care of the bill (*known as mobile-commerce or m-commerce*).

Whilst this initial use may seem rather mundane, it must be noted that this technology has the potential to have a significant impact on Internet systems. For example, this system could be incorporated with the Screenfridge listed above, making it possible for the fridge to detect what items are within (*and possibly in nearby cupboards*), and could automatically connect onto an Internet database to determine the recipes that the user can follow with the available items.

1.5.3 Mobile Communication Devices

Mobile communication companies have largely concluded that there will be three broad categories of mobile communicators. Mitra (1997) identifies the following categories and their perceived associated technologies:

- Communicator
  - Voice, Data, Fax
  - Smart Messaging
  - pJAVA
- Smart Phone
  - Voice
  - Smart Messaging
  - eJAVA
- Cellular Phone
  - Voice
  - Smart Messaging

It should be noted that Mitra (1997) is a publication by Sun Microsystems, the proprietors of Java. However, current trends do suggest that Java will play an important part in future systems because it provides a cross platform programming standard (*although not currently open source*). One example of its success is the amount of Java mobile phone games from companies such as MyPhoneGames.co.uk (2004).

<u>OQO model 01 - uPC</u>

Figure 1.3
Source: Zucker (2004, p.16)

It is evident that the mobile devices illustrate different degrees of convergent technology. The most extreme is the Communicator that fully combines PDA and mobile phone systems. This device has the potential to be fully integrated with Internet systems, whilst the Cellular Phone omits this functionality and provides similar features to current mobile phones.

This functionality has largely been achieved by the mobile communication companies. For example, Hutchison 3G UK Limited offer a number of facilities which include video calling (*in addition to voice calling*), location services, multimedia messaging, e-mail/fax services, entertainment (*games, movie clips, sport videos, horoscopes etc.*), news and finance data repositories, travel information, are examples of the many services provided (Hutchison (2004)). These are available on a range of mobile handsets which can be categorised in accordance with Mitra (1997).

However, there appears to be some evidence which suggests there is demand for traditional PC functionality on PDA-sized mobile device which could converge with mobile phone technology. 4G.co.uk (2004) state that mobile phone operators "*have demanded products that will offer PC capabilities in a PDA form factor*".

Figure 1.3 provides an illustration of the OQO model 01 ultra personal computer (uPC). This device provides a fully-functional Windows XP computer that is merely 4.9 inches long, 9 inches wide, 0.34 inches thin, and weighs only 14 ounces. Zucker (2004, p.16) states that he believes the principle usage of the product will be as a primary desktop computer, capable of connecting to fully sized external peripherals, such as keyboard, monitor, and disk drives. We note that Zucker (2004) does not discuss the use of this device with telecommunications.

A further interesting and related technology is the development of projection keyboards. Canesta (2004) claim to offer "*the world's first projection keyboard capable of being fully integrated by OEMs into smart phones, cell phones, PDAs, or other mobile or wireless devices*". This device is capable of projecting the image of a full-sized keyboard onto a convenient flat surface between the device and the user, such as a tabletop, and enables a user to press virtual buttons. See Figure 1.4.

Canesta projection keyboard

Figure 1.4
Source: Canesta (2004)

It has been evident from general reading that the categories that Mitra (1997) described are well established. However, the next generation of mobile device is very difficult to predict due to the number of constraints associated with this technology. In spite of this, it is clear that many companies are actively involved in overcoming usability and functionality constraints to provide the user with increased ability.

1.5.4 Sensors

The word 'sensor' covers a vast range of devices. Indeed, all of the devices discussed so far can be considered sensors. However, there are a significant number of authors that predict a "sensor boom" where a vast number of tiny devices, such as thermostats, lighting controls, environmental sensors, and medical monitors are implemented. Hardy (2004) states that networks consisting of these automated devices will be capable of generating tremendous amounts of information, allowing scientists to discover patterns that provide insight into "how the world works".

Low Power Sensors

ZigBee Alliance (2004a) describe themselves as "*an association of companies working together to enable reliable, cost-effective, low-power, wirelessly networked, monitoring and control products based on an open global standard*". The consortium consists of leading semiconductor manufacturers, technology providers, OEMs, and end users that all have a shared interest in this technology. Brands include: Honeywell, Invensys, Mitsubishi Electric, Motorola, Philips, and Samsung.

Summary of 802.15.4:
> low-data-rate
> low-power-consumption
> low-cost applications

The ZigBee consortium is focused on the IEEE 802.15.4 standard for low-rate Wireless Personal Area Networks (WPANs). This standard must be distinguished from its IEEE 802.15.3a counterpart (*also known as ultra wideband or UWB*), discussed later in this section.

Zheng and Lee (2004) provide various examples to describe where and how this technology could be used.

1) Life saving – the sensors could be used to detect and shutdown gas leakage.
2) Household Convenience – the sensors could be used to mute the television when the phone rings, or dim a room's lights when the TV is turned on.
   Furthermore, Zheng and Lee describe prioritised household user profiles, where members of a house have an 802.15.4 device and a corresponding private electronic profile capable of adjusting the lights, temperature, music, and TV channel etc.
3) Enhanced GPS – the sensors could be combined with signposts and other displays to provide more detailed and accurate information to motorists. They could provide the traffic status, whether a road is one or two way, and offer relevant accident information. Similarly, this technology could enhance public transport by providing detailed vehicle location details and the corresponding estimated time of arrival.

However, they are quick to state that these particular applications may not become reality. Furthermore, they state that it may take some time to prove the technology and create a viable market. Figure 1.5 puts ZibBee into context with other wireless technologies.

Diagram Illustrating various wireless technologies in terms of bandwidth and range

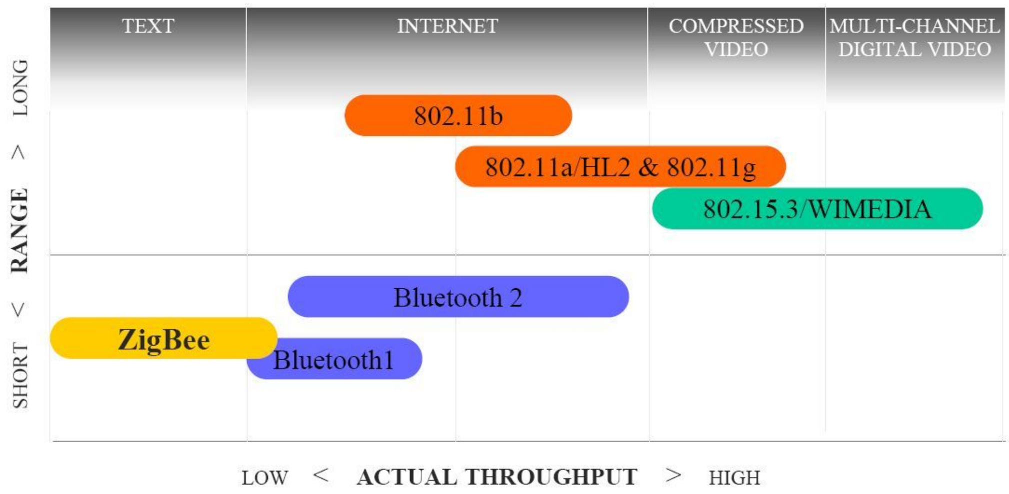

Figure 1.5

Source: Heile (2003)

High Power Sensors

This topic is potentially very large and is a research area of its own. Consequently, for the purposes of this document UWB is the only standard to be discussed.

Ultrawidebandplanet.com (2004) state that UWB is a wireless communications technology capable of transmitting data in short pulses which are spread over a relatively wide section of the electromagnetic spectrum. There are two major advantages over single frequency transmissions:

1) Data can be transmitted in large bursts (*high bandwidth*) because the data is spread across several frequencies simultaneously. One's initial reaction may be that this would interfere with other communications systems, however because UWB devices only transmit for very short periods, there isn't any apparent interference with other transmissions.

Consequently
2) Frequencies employed by other systems can be utilised because UWB transmissions merely appear as white noise, thus eliminating the interference concerns.

In spite of these benefits, the implementation of this technology has largely been stifled by the various radio spectrum regulators. These have expressed grave concern about interference with existing radio transmissions, such as those which Global Positioning Systems (GPSs) use.

However, Freescale Semiconductor Inc. claimed, on the 9th August 2004, to be the first company to receive Federal Communications Commission (FCC) certification for its Ultra-Wideband (UWB) communications solution; see Electronic News (2004) for further details.

Other countries, such as UK, have enabled UWB solutions before 2004, see Roke Manor Research (2002, p.11), however the FCC approval is significant for widespread adoption of the technology.

Devices that employ this technology are able to transmit data through walls, and communicate at up to one Gigabit per second. This makes it a suitable technology for wireless Local Area Networks (WLANs) and Personal Area Networks (PANS), especially for domestic entertainment, where it could be possible to transmit several videos streams simultaneously; for example, connecting a DVD player or surveillance system to the TV without the need for cables.

It is assumed that such devices could connect to the Internet via a wireless Internet Access Point.

Bibliography

Abilene (2003) *IPv6,* Michigan, Internet2.
Available from: http://abilene.internet2.edu/services/ipv6.html, [Date Accessed: 21 November 2003]

Apple Computer Inc. (2004) *Apple-iPod*, California, Apple Computer Inc.
Available from: http://www.apple.com/ipod/, [Date Accessed: 23 June 2004]

BBC (2004) *Europe launch for Apple's iTunes*, London, British Broadcasting Corporation.
Available from: http://news.bbc.co.uk/2/hi/entertainment/3805565.stm, [Date Accessed: 23 June 2004]

Bocij, P., Chaffey, D., Greasley, A., Hickie, S. (2003) *Business Information Systems. Technology Development and Management* Second edition, London: Pearson Education Limited.

Canesta (2004) *Canesta, Inc. -- Products*, California, Canesta.
Available from: http://www.canesta.com/products.htm, [Date Accessed: 20 September 2004]

ChildLocate (2004) *Childlocate-Mobile Phone Tracking*,Guildford UK, ChildLocate.
Available from: http://www.childlocate.co.uk/, [Date Accessed: 24 June 2004]

Department of Electrical Engineering and Computer Sciences (1999) Mark D. Weiser Bio, California, Berkeley.
Available from: http://www.cs.berkeley.edu/Weiser/bio.shtml, [Date Accessed: 24 June 2004]

Dertouzos, M. L. (2001) *Future computing: Human-centric systems*, Massachusetts, Accela Communications, Inc.
Available from: http://utilitycomputing.itworld.com/4605/CWD010507STO60203/page_1.html, [Date Accessed: 27 June 2004]

Digi-Frame Inc. (2004) *::: Digi-Frame Inc. - Digital Picture Frames :::*, New York, Pragmatic Designs Inc.
Available from: http://www.digi-frame.com/home.html, [Date Accessed: 24 June 2004]

Dingli, A. (2001) *Internet Enabled Devices*, Sheffield, The University of Sheffield.
Available from: http://www.dcs.shef.ac.uk/~alexiei/WebSite/University/MultiPLAT/WebDocs/AdditionalDocs/Internet%20Enabled%20Devices.htm, [Date Accessed: 22 June 2004]

Electrolux (2004) *Electrolux* – Screenfridge, Stockholm, AB Electrolux.
Available from: http://www.electrolux.com/screenfridge/, [Date Accessed: 23 June 2004]

Electronic News (2004) *Electronic News - Freescale Lands First FCC Certification for UWB*, Ohio, Reed Elsevier Inc.
Available from: http://www.reed-electronics.com/electronicnews/article/CA443752, [Date Accessed: 25 August 2004]

Electronic News (2002) *Is Wider Really Better?*, Ohio, Reed Elsevier Inc.
Available from: http://www.reed-electronics.com/electronicnews/article/CA240098, [Date Accessed: 25 August 2004]

Foroohar, R. (2004) *The Future of Shopping*, New York, Newsweek, Inc.
Available from: http://msnbc.msn.com/id/5093197/site/newsweek/, [Date Accessed: 26 June 2004]

Hardy, Q. (2004) *Forbes.com: The Wireless Boom That Won't Need Us*, California, Data Depth Corporation.
Available from: http://www.forbes.com/technology/networks/2004/06/09/cz_qh_0609robonet.html
[Date Accessed: 24 June 2004]

Heile, B. (2003) *ZigBee Alliance Overview, Goals and Objectives*, California, ZigBee Alliance.
Available from: http://www.zigbee.org/events/documents/November2003_Open_House_Presentations/

Other/02BobHeile.pdf, [Date Accessed: 24 August 2004]

Hornby, R. M (1999) RFID Solutions for the express parcel and airline baggage industry,*RFID Technology* (Ref. No. 1999/123), IEE Colloquium.

Howarth, R. (2004) Personal Communication 23 June 2004.

Hutchison 3G UK Limited (2004) *Services: Overview*, Hemel Hempstead, Hutchison 3G UK Limited.
Available from: http://three.co.uk/explore/services3/detail.omp, [Date Accessed: 19 September 2004]

IEEE Computer Society (1998) *Virtual Roundtable: Ubiquitous Computing*, Washington, Institute of Electrical and Electronics Engineers, Inc.
Available from: http://www.computer.org/internet/v2n2/round1.htm, [Date Accessed: 24 June 2004]

Mitra, S. (1997) *Personal Communicators of the Future*, California, Sun Microsystems Inc.
Available from: http://java.sun.com/javaone/sessions/slides/JC04/JC04_4.ppt.pdf, [Date Accessed: 27 June 2004]

MyPhoneGames.co.uk (2004) *Mobile Phone Games, Java Games, Mobile Games*, Kentucky, MaximumASP LLC.
Available from: http://www.myphonegames.co.uk/, [Date Accessed: 27 June 2004]

Napster (2004) *Napster.co.uk*, California, Roxio Inc.
Available from: http://napster.co.uk/, [Date Accessed: 24 June 2004]

NTL (2004) *ntl: - Simply Better Digital TV -ntl vs Sky, Newport, NTL Group Ltd*
Available from: http://www.ntlhome.com/ntl_tv/sky.html, [Date Accessed: 22 June 2004]

NTL (2002) *Text on TV*, Newport, NTL Internet Ltd.
Available from: http://www.digitalcabletv.co.uk/doc_disp.asp?HTDocID=479, [Date Accessed: 22 June 2004]

PalmOne Inc. (2004) *palmOne - Software - Palm Web Browser Pro*, California, Palm Inc.
Available from: http://www.palmone.com/us/software/palm/, [Date Accessed: 22 June 2004]

Picsel Technologies (2002) *Browser Product Information Sheet Picsel*, Glasgow, Picsel Technologies Ltd.
Available from: http://www.picsel.com/pdf/picselbrowser.pdf, [Date Accessed: 27 June 2004]

Roke Manor Research (2002) *Sensor Technology*, Hampshire, Roke Manor Research Limited.
Available from: http://www.roke.co.uk/download/brochures/sensor_technology.pdf, [Date Accessed: 25 August 2004]

Schwarz, A. (2001) *ISRC Future Technology Topic Brief - Human-Centric Computing*, Texas, University of Houston.
Available from: http://www.uhisrc.com/FTB/Humancentric/Humancentric.pdf, [Date Accessed: 25 June 2004]

ScienceDirect (2004) *ScienceDirect*,Elsevier B.V., Amsterdam.
Available from: http://www.sciencedirect.com/, [Date Accessed: 24 June 2004]

Snowden, O.F.J. (2004) *Entity Location Service (ELS) Phase 1 Specification IPv6 (Internet Protocol version 6) Address Translation specific to domain of interest*, Southampton: University of Southampton.

Snowden, N. (2004), Personal Communication 22 June 2004.

Tanenbaum, A. S. (2003), *Computer Networks* Fourth edition, New Jersey: Prentice Hall.

Thompson, K. (2000) *Internet Appliances*, California, Rogers Communications Inc.
Available from: http://www.kristi.ca/web_appliances/content.html, [Date Accessed: 23 June 2004]

Ultrawidebandplanet.com (2004) *Ultrawidebandplanet.com - FAQ:Frequently Asked Questions (FAQ): Ultra*

Wideband (UWB), Connecticut, Jupitermedia Corporation.
Available from: http://www.ultrawidebandplanet.com/faq/article, [Date Accessed: 25 August 2004]

Whatis.com (2004) *Internet appliance*, Massachusetts, TechTarget Inc.
Available from: http://whatis.techtarget.com/definition/0,,sid9_gci914561,00.html, [Date Accessed: 22 June 2004]

Xdrive (2004) *Xdrive | Secure Online Storage*, California, Xdrive Inc.
Available from: http://www.xdrive.com/, [Date Accessed: 24 June 2004]

Weiser, M. (1996a) *Open House*, California, Xerox PARC.
Available from: http://www.ubiq.com/hypertext/weiser/wholehouse.doc, [Date Accessed: 24 June 2004]

Weiser, M. (1996b) *Ubiquitous Computing*, California, Xerox PARC.
Available from: http://www.ubiq.com/hypertext/weiser/UbiHome.html, [Date Accessed: 24 June 2004]

Weiser, M. (1996c) IN Department of Electrical Engineering and Computer Sciences(1999)

ZigBee Alliance (2004a) *About the Alliance*, California, ZigBee Alliance.
Available from: http://www.zigbee.org/about/index.asp, [Date Accessed: 24 August 2004]

Zheng, J. and Lee, M. J. (2004) Will IEEE 802.15.4 make ubiquitous networking a reality?: a discussion on a potential low power, low bit rate standard, *IEEE Communications Magazine*, 42(6), 140-146.

Zhuge, H. (2002) Clustering Soft-devices in the Semantic Grid, *IEEE Computing in Science & Engineering*, 4(6), 60-62.

Zucker, D. F. (2004) New Products: OQO POCKET-SIZED PC, *IEEE Pervasive Computing*, 3(3), 16-17.

4G.co.uk (2004) *4G Mobile PC Capabilities*, Newport, 4Gcouk Limited.
Available from: http://www.4g.co.uk/PR2004/March2004/2015.htm, [Date Accessed: 19 September 2004]

Section 2: Analysis of current and planned structure of the Internet

2.1 Physical Structure of the Internet today

Woodcock (2004) states that the "*Internet is an internetwork, a global conglomeration of smaller networks able to communicate and transfer information among themselves*". Figure 2.1 provides an overview of the Internet today.

Physical Structure

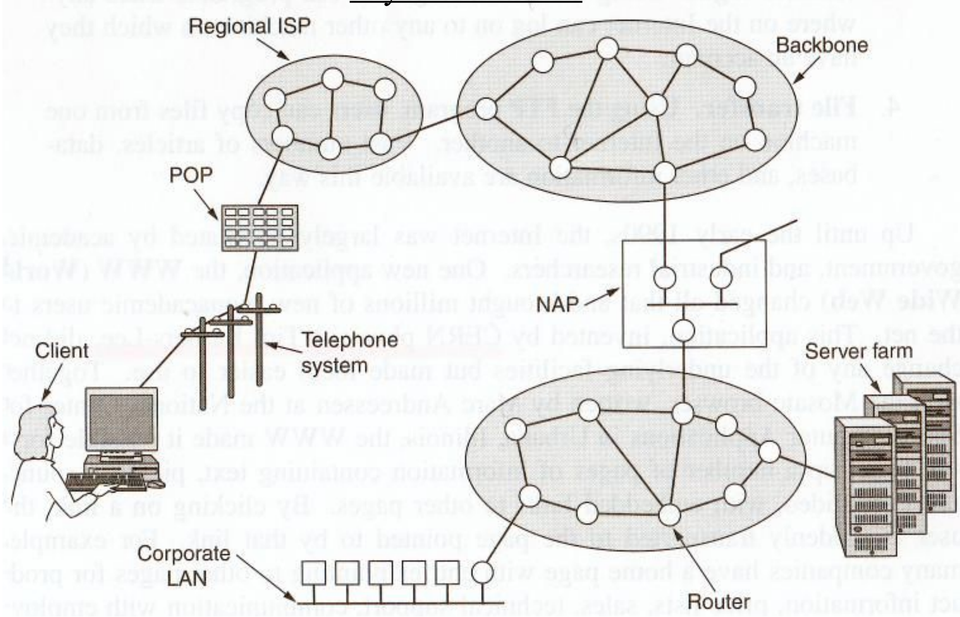

Figure 2.1 Source: Tanenbaum (2003, p.58)

- Clients (*users*) can connect to their Internet Service Provider (ISP) via a range of means:
  - a standard phone line
  - Digital Subscriber Line (DSL)
  - Cable
  - Satellite

    A modem is typically required to access a communications medium. They then establish a point of presence (POP).
- The ISP's regional network consists of interconnected routers, enabling communication between nodes within the ISP. However, the ISP is connected to a "backbone" operator.
- The backbone contains thousands of routers, typically connected by high-bandwidth fibre optics, which provide access to the global Internet. They are operated by the telecom conglomerates such as AT&T, Sprint and BT. See Figure 2.2 for the backbone structure of the UK and the rest of Europe.
  1) there are a number of "backbones" throughout the world, which are connected to one another via Network Access Points (NAPs).
  2) within the UK there are several backbones including UUNET, BTnet and JANET (*UK academic and research network*). This is typically the case in other parts of the world, where several companies have their own backbone in a geographical region.

Graphical Representations of a UK and European backbone

UUNET'sUK Backbone in 1999

Source: Dodge (2004)

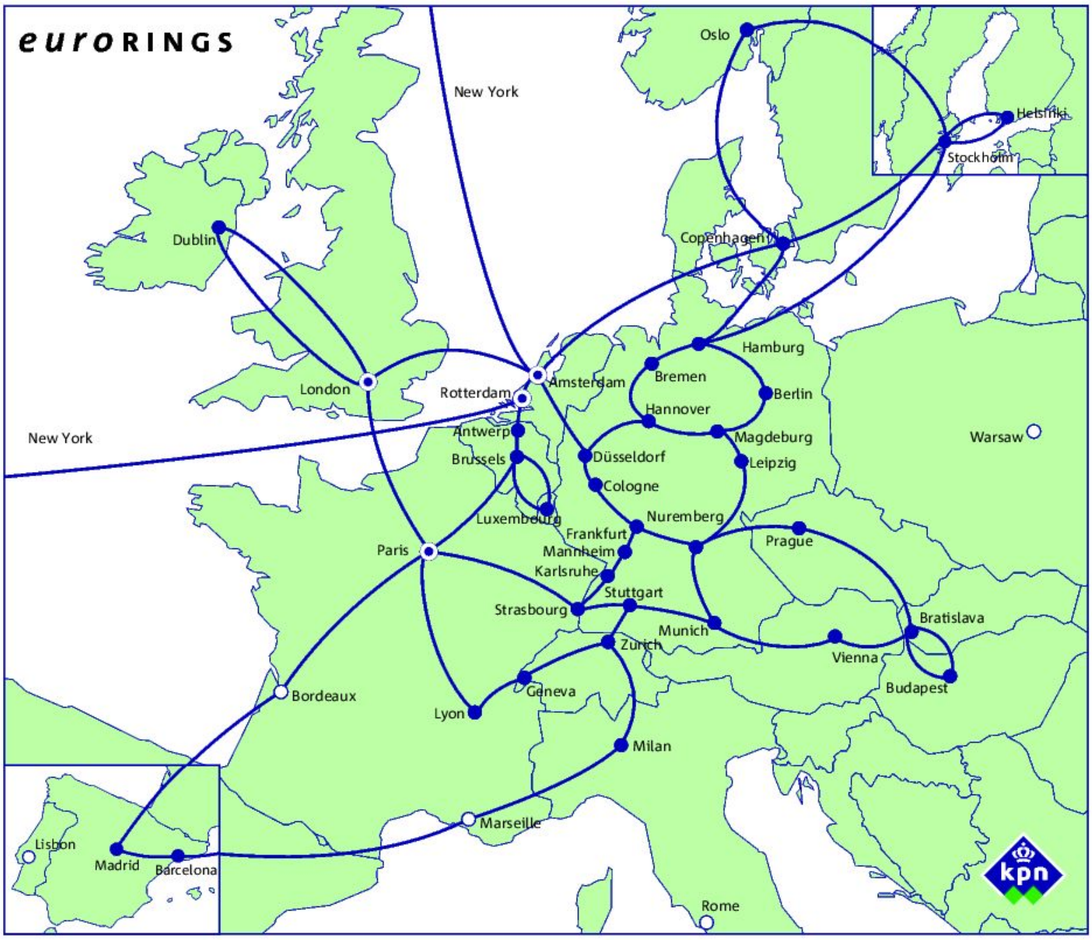

EuroRINGS European Backbone in 2004

Figure 2.2 Source: EuroRINGS International Services (2004)

- Server farms (*machines serving thousands of Web pages per second*), are connected directly onto the Internet backbone for performance reasons. Corporate LANs are typically integrated closely with server farms, although smaller organisations may be connected via a standard ISP. (Explanation summarised from: Tanenbaum (2003, p.58))

2.2 Virtual Structure of the Internet today

The virtual structure is important to the Internet community because essentially it defines the conceptual model of the Internet.

Two fundamental principles of the Internet's Virtual Structure are, according to the American Registry for Internet Numbers (2004):

- ***Internet Protocol (IP) Address***

  The IP address defines the identity of an Internet connected device, and the IP dictates how information passes the many networks, Bisson (2000). IP addresses enable routing of packets (*chunks of data*) throughout the world, and the IP address ranges are controlled by four organisations:

  ARIN – American Registry for Internet Numbers – *American Registry*
  RIPE NCC – Réseaux IP Européens Network Co-ordination Centre – *European Registry*
  APNIC – Asia Pacific Network Information Centre – *Asia Pacific Registry*
  LACNIC – Latin American and Caribbean Network Information Centre – *Latin American and Caribbean Registry*

  However, IP addresses are difficult for humans to recognise and remember. Consequently, a naming system was required to represent the IP addresses. Initially this was a simple text file (*hosts.txt*) linking names to IP addresses, although this solution wasn't scalable which gave rise to the DNS (*see below*).

- ***Domain Name System (DNS)***

  DNS is a classification system, devised by Mockapetris (1987), to provide a consistent name space for referring to resources. Further to this, DNS distributes the processing required to translate node names into IP addresses which makes the system scalable.

  DNS has given the Internet a logical structure (*despite it physically being chaotic*) and allows for an orderly and continual growth.

  An example of how DNS works:

> E.g.
> Assume the site with Universal Resource Locator (URL) ns.somewhere.com. needs to be contacted and no address caching is used. The address is traversed from right to left.

- Initially, one of the root "." domain name servers (*which know the top level domain names such as .com, .au and .uk etc.*) is contacted and returns the IP address of the ".com" DNS server(s).
- The ".com" is then queried and returns the IP address of ".somewhere".
- ".somewhere" is finally contacted which knows the address of "ns" and returns its IP.

The processing and data required for domain IP lookup is distributed between nodes.

2.3 Internet Protocols (Client Level)

Internet protocols are vitally important because they specify the communication standard that Internet devices use to communicate.

It is convention to communicate using a protocol on a protocol specific **port**, an access means to a computer system. IANA (2004) provide a full list of known ports and their intended service.

Note: a port is merely an access means to a computer system. Alone, they don't specify the protocol, although there are conventions which should be adhered to, see IANA(2004). For example, port 80 is the standard hypertext transfer protocol (HTTP) access port).

2.3.1 Example specification of Internet Protocols in a standard web-browser.

The following is a valid string for connecting to a web page:
- "http://www.somewhere.com.:80"

One can dissect this as:

- "http:" - This informs the web-browser that it is to expect data conforming to the HTTP protocol.
- "//" - The double slash (//) informs the web-browser that the next part is the name of the server.
- "www.somewhere.com." - this a valid Domain Name to obtain its corresponding IP address.
- ":80" - this informs the web-browser that it is to communicate with the server on port 80.

Typically, the default protocol in a web-browser is "http", which it will (*typically*) automatically associate with port 80 without the user specifying otherwise. Other protocols are: ftp, gopher, https, telnet, etc.

2.4 The Planned Physical Internet Structure

Research suggests that the Internet will continue to evolve, demanding more advanced and faster resources. TeleGeography Research (2003) state that (based on data from the first quarter of 2003) "*international Internet traffic is growing at an annual rate of 67 percent*". Figure 2.3 provides the percentage bandwidth growth between the years 2000-2003; it is evident that the world's annual bandwidth growth is significant.

This is ultimately explained by the increase in Internet user demands. Sullivan (2004) states that one-third of all adult Americans have a broadband connection at home or work, and that broadband growth is especially strong with home users. Home user Internet activities normally include web-browsing, instant messaging, voice over IP telephony, game playing, software updates, and listening/watching to streaming audio/video. Many of these services have different requirements and are putting an increasing demand on the current physical hardware. Some services require real-time and relatively small bandwidth networks (*e.g. Instant messaging*), whilst others simply require large amounts of bandwidth (*e.g. File downloading*).

International Internet Bandwidth Growth by Region, 2000-2003

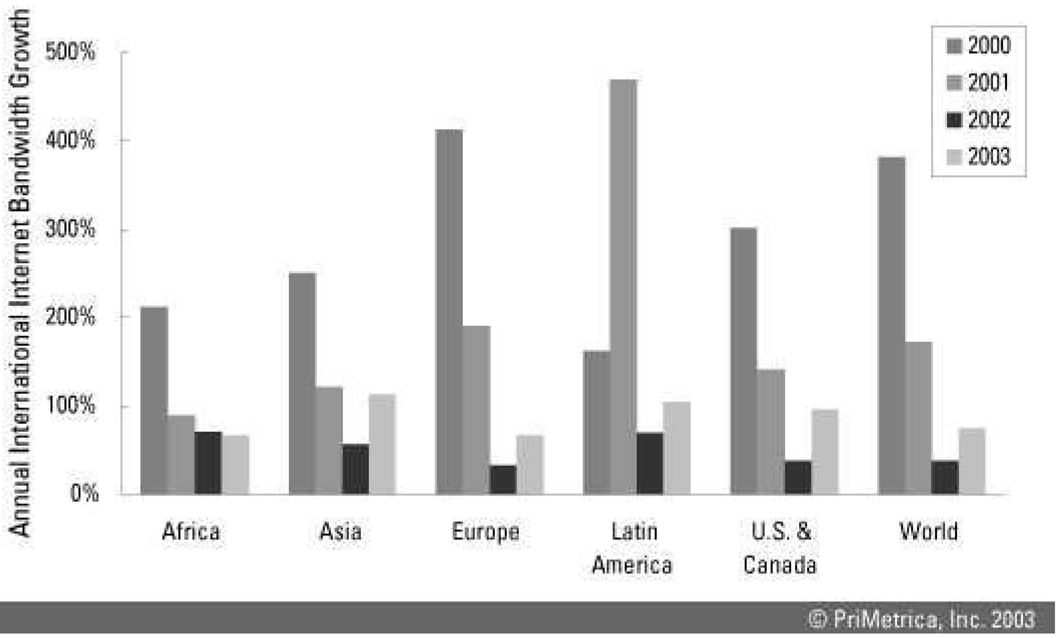

Figure 2.3 Source: TeleGeography Research (2003)
Note: Permission was granted to use image.

Further to simply increasing network capacity, the physical network structure of the Internet is changing to accommodate the new and evolving IP services. Cisco Systems (2000) state that organisations are looking for strategies to integrate disparate networking into one single common network infrastructure. Volpi (2004) reiterates this and states that for the past hundred years every time an organisation needed to implement a new application, a new network had to be built. However, on 24 May 2004 Cisco Systems revealed their new technology based on the fundamental principles that organisations only want (*primarily due to cost*) one network infrastructure, that will be future proof and compatible with the past. Gofus (2004), for the MCI telecommunications company, states that it is the ability to differentiate service and support a variety of digital media over the IP infrastructure that is vital.

Grubb (2004) provides a demonstration of a 40Gbps system (single CRS-1 nodes) able to support the following services simultaneously:

- run a high-definition live video conference with a MCI representative several miles away,
- simulated 4000 MP3 downloaders (*including 5 real users to confirm QoS*),
- simulated 100,000-125,000 gamers on-line (*including some real users to confirm QoS*),
- simulated 2500 IP set-top boxes + 200 televisions displayed as part of the demonstration,

- protocol installation to support IP videophones (*whilst maintaining ALL network activity – QoS was maintained*),
- simulated 1000 videophone conversations + several customer videophone conversations as part of the demonstration to confirm QoS,
- simulated 1,000,000 web-browser users.

It is evident that the above demonstration referred to a specific product, yet it demonstrated that future Internet technology will be required and able to support users with different QoS needs simultaneously across the same communications medium.

Trusted computing (TC) is a further aspect that in the long term may be adopted within the Internet structure. The Trusted Computing Group (2003, p.2) state that "*concerns about the security of communications, transactions, and wireless networks are inhibiting realization of benefits associated with pervasive connectivity and electronic commerce*". The group's plans are to make computing (*and therefore the Internet*) more secure by working with both hardware and software vendors. The infrastructure could be used to enforce digital rights management (DRM), ensuring users only have access to "legal" software and entertainment and enable authorities to carry out what many describe as "Big Brother" activities. This could split the Internet into two distinguishable structures and categorise users into those prepared to fully accept TC, those that totally reject TC, and those that use both systems (*e.g. TC for on-line banking etc. and non-TC for other activities such as movie watching*).

An example of a TC failure was the introduction by Intel of unique CPU identifiers. Critics complained that the technology could expose computer users to privacy intrusions. This prompted such strong public opposition, with many customers putting a boycott on Intel products, that Intel swiftly disabled this functionality, citing there must be an acceptable balance between privacy and security. Source: O'Harrow, R. and Corcoran, E. (1999)

It is evident from the convergence of networks, TC and the new wireless standards (*e.g. 802.11 and 802.16*), that the entire Internet is becoming more mature and wider reaching, essentially laying the foundations for the evolving software applications.

2.5 Planned Virtual Structure

The planned virtual structure of the Internet appears to be uncertain. IPv6 (Internet Protocol version 6) was planned to be the successor to IPv4 for a number of reasons, summarised from: Tanenbaum (2003, p.472):

- It provides a very large address space (2^{128} *addresses*), thus enabling more devices to connect to the Internet simultaneously;
- is more efficient than IPv4 (*the protocol structure and routing tables are simplified*);
- and it provides better support for QoS.

Note: features of IPv6 which have since been adopted in IPv4 are ignored and only the most significant reasons have been listed.

Further to IPv6, much work has taken place in Mobile IPv6. This is a protocol which enables nodes to remain reachable whilst roaming in the IPv6 Internet and can be likened to cellular phone roaming, where a mobile phone is reachable (*by the same phone number*) in different locations.

However, Francis (2003, p.1) sheds doubt over the entire IPv6 concept and states that IPv6 may never be widely adopted. He was the founder of Network Address Translation (NAT) which he created in 1993. At that time, he thought it was a temporary solution giving organisations extra time before the IPv4 addresses were depleted, but he now sees it as part of a technology that relieves the need for IPv6.

He cites: "*If the commercial world hasn't accepted IPv6 so far, what will make it do so in the future?... Nothing...NAT provides enough addresses for each human on the planet to have an endless number of private connections and about 250 simultaneous active 'global' connections*."

His proposed alternative to IPv6, NUTSS (*see below*), accomplishes the same end result, although in a far more complex way:

- **N**AT effectively extends the IP address space.
- **U**RIs (Uniform Resource Identifier) restore end-to-end stable addressing; these look similar to e-mail addresses, such as user@domain.com or process@domain.com.
- **T**unnels allow protocols like IPsec and mobile IP (*and even TCP*) to run through NATs over UDP (Universal Datagram Protocol).
- **S**IP (Session Initiation Protocol) routes messages with URIs, end-to-end, and lets hosts signal their intentions to each other in real time.
- **S**TUN (simple traversal of UDP through NATs) tells hosts how to establish direct IP connectivity through NATs.

However, in spite of the Francis' view, the general consensus is that IPv6 will replace IPv4. Koprowski (2003) states that the governments of Japan, European Union, and United States have mandated that, in the near future, network devices must support IPv6.

For "*Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)*", see Austein (2002).

Bibliography

Austein, R. (2002) *Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)*, Virginia, The Internet Society.
Available from: ftp://ftp.rfc-editor.org/in-notes/rfc3364.txt, [Date Accessed: 2 July 2004]

American Registry for Internet Numbers (2004) *Registration Services*,
Available from: http://www.arin.net/registration/index.html, [Date Accessed: 30 June 2004]

Bechtold, S. (2003) *The Present and Future of Digital Rights Management – Musings on Emerging Legal Problems*, Tübingen, Eberhard Karls Universität Tübingen.
Available from: http://www.jura.uni-tuebingen.de/bechtold/pub/2003/Future_DRM.pdf, [Date Accessed: 2 July 2004]

Bisson, S. (2000) *How the Internet works: IP*, Bath, PC Plus.
Available from: http://www.pcplus.co.uk/tutorials/default.asp?pagetypeid=2&articleid=4430&subsectionid=376, [Date Accessed: 30 June 2004]

Cisco Carrier Routing System Enabling Next-Gen IP Networks (2004) Webcast, Conducted by Cisco Systems. California: Cisco Technology Inc.
Available from: http://cisconews.feedroom.com/index.jsp?fr_story=6c7d28661aeeafa22a0f134f0a35d8eeed487c8b, [Date Accessed: 2 July 2004]

Cisco Systems (2000) *Technical Considerations for Converging Data, Voice, and Video Networks*, California, Cisco Technology Inc.
Available from: http://www.cisco.com/warp/public/cc/so/neso/vvda/avvid/tecon_wp.htm, [Date Accessed: 2 July 2004]

Cisco Systems (2004a) *Cisco - Multiservice Networking*, California, Cisco Technology Inc.
Available from: http://www.cisco.com/warp/public/cc/so/neso/vvda/avvid/, [Date Accessed: 2 July 2004]

Deering, S. and Hinden, R. (1998) *Internet Protocol, Version 6 (IPv6) Specification*, Virginia, The Internet Society.
Available from: http://www.faqs.org/rfcs/rfc2460.html, [Date Accessed: 1 July 2004]

Dodge, M. (2004) *An Atlas of Cyberspaces- ISP Backbone Maps*,
Available from: http://www.cybergeography.org/atlas/isp_maps.html, [Date Accessed: 29 June 2004]

EuroRINGS International Services (2004) *KPN – EuroRings network*, The Netherlands, KPN International.
Available from: http://www.kpn.com/kpn/show/id=397137/sc=3b4042, [Date Accessed: 29 June 2004]

Fairlie, R. (2003) *In matters of privacy, whom can you trust?*, California, CNET Networks Inc.
Available from: http://reviews.cnet.com/4520-3000_7-1009386-1.html, [Date Accessed: 2 July 2004]

FAQS.ORG (2004) *RFC By Category*, Illinois, FAQS.ORG.
Available from: http://www.faqs.org/rfcs/np.html, [Date Accessed: 30 June 2004]

Francis, P. (2003) Is the Internet Going NUTSS?, *IEEE Internet Computing*, 7 (6), 94-96.

Freesoft (1997) *HTTP Protocol Overview*, Maryland, Freesoft.
Available from: http://www.freesoft.org/CIE/Topics/102.htm, [Date Accessed: 2 July 2004]

Graham, I. (1998) *URLs for HTTP Servers*, Toronto, University of Toronto.
Available from: http://www.utoronto.ca/webdocs/HTMLdocs/NewHTML/url_http.html, [Date Accessed: 1 July 2004]

Gofus, N. (2004) IN *One Network for the Future* (2004).

Grubb, J. (2004) IN *Cisco Carrier Routing System Enabling Next-Gen IP Networks* (2004).

Howarthson, T. (2004) *BT to save £1bn a year with IP network*, London, Situation Publishing.
Available from: http://www.theregister.co.uk/2004/06/09/bt_ip_network/, [Date Accessed: 2 July 2004]

IANA (2004) *Port Numbers*, California, The Internet Assigned Numbers Authority.
Available from: http://www.iana.org/assignments/port-numbers, [Date Accessed: 30 June 2004]

Interoute (2004) *Glossary*, Interoute Telecommunications (UK) Ltd, London.
Available from: http://www.interoute.com/glossary.html, [Date Accessed: 28 June 2004]

Koprowski, G. J. (2003) *Internet Protocol for the Future: IPv6 Poised for Adoption*, California, ECT News Network Inc.
Available from: http://www.technewsworld.com/story/31230.html, [Date Accessed: 2 July 2004]

Mockapetris, P. (1987) *DOMAIN NAMES - CONCEPTS AND FACILITIES*, Virginia, The Internet Society.
Available from: ftp://ftp.rfc-editor.org/in-notes/rfc1034.txt, [Date Accessed: 30 June 2004]

O'Harrow, R. and Corcoran, E. (1999) *Intel Drops Plans to Activate Chip IDs*, Washington DC, The Washington Post Company.
Available from: http://www.washingtonpost.com/wp-srv/washtech/daily/jan99/intel26.htm, [Date Accessed: 2 July 2004]

One Network for the Future (2004) Webcast. Conducted by Cisco Systems. California: Cisco Technology Inc.
Available from:
http://cisconews.feedroom.com/index.jsp?fr_story=5414aec62185948a1aed2074b896b407da185b2b, [Date Accessed: 2 July 2004]

Sullivan, A. (2004) *Net users demand wider broadband*, New Delhi, Cyber India Online Ltd
Available from: http://www.ciol.com/content/search/showarticle1.asp?artid=56536, [Date Accessed: 2 July 2004]

Tanenbaum, A. S. (2003) *Computer Networks* Fourth edition, New Jersey: Prentice Hall.

TeleGeography Research (2003) *Executive Summary: Global Internet Geography*, Exeter, PriMetrica Inc.
Available from: http://www.telegeography.com/products/gig/exec_sum.php, [Date Accessed: 2 July 2004]

Trusted Computing Group (2003) *Truster Computing Group Backgrounder*, Oregon, Trusted Computing Group.
Available from: https://www.trustedcomputinggroup.org/downloads/TCG_Backgrounder.pdf, [Date Accessed: 2 July 2004]

Valente, O. (2004) IN *One Network for the Future* (2004).

Volpi, M. (2004) IN *One Network for the Future* (2004).

Woodcock, J. (2004) *Microsoft TechNet: Structure of the Internet and the Web*, California, Microsoft Press.
Available from: http://www.microsoft.com/technet/itsolutions/network/evaluate/technol/intweb.mspx, [Date Accessed: 28 June 2004]

Section 3: Investigation of the Internet as an Object Oriented distributed database

3.1 Definition of an Object Oriented (OO) database

An OO database provides the persistent storage of objects and facilitates all the aspects of the OO data model (*discussed below*).

Note: an OO database inherits all the features of a standard database such as Data Definition, Data Maintenance, Data Retrieval and Data Control. See Beynon-Davies (2004, p.36).

Thomson Learning (2004) confirms this and states that an OO Database "*keeps track of objects, entities that contain both data and the action that can be taken on the data*". They further note that this facilitates the storage of non-traditional data (*i.e. non text*) such as photographs, video, and audio clips.

3.1.1 Objects

Beynon-Davis (2004, p.116) defines an object as a package of data and procedures, where data are contained by an object's attributes and the procedures (*actions that can be performed*) defined by an object's methods.

Figure 3.1 provides a visual representation of an object. The single and double headed arrows represent one-way and two-way communication respectively (*i.e. Procedures can communicate without being invoked*).

Graphical Representation of an Object

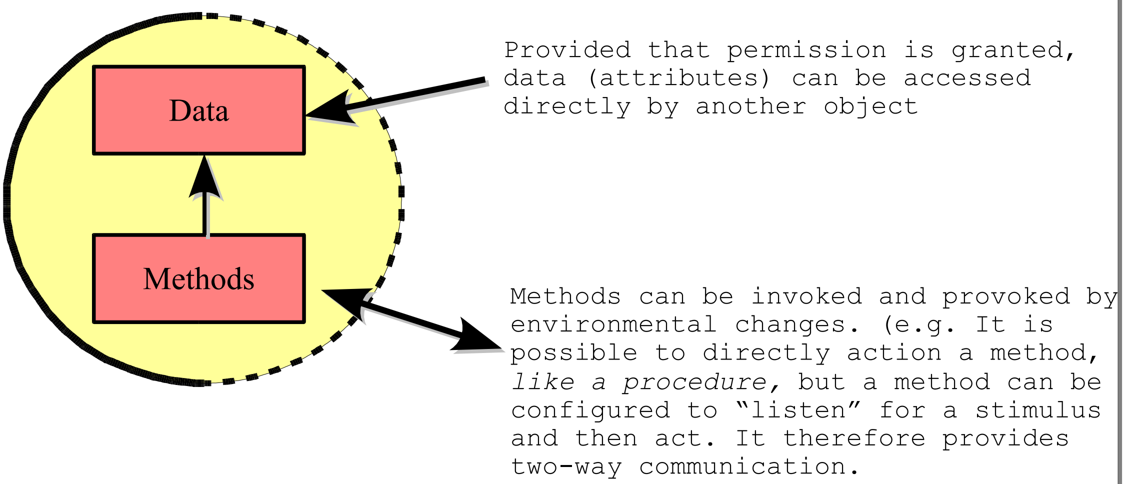

Figure 3.1

Figure 3.2 provides an overview of inter-object communication. Sun Microsystems (2004a) state that objects interact and communicate by sending messages to each other.

Graphical Representation of communication between Objects

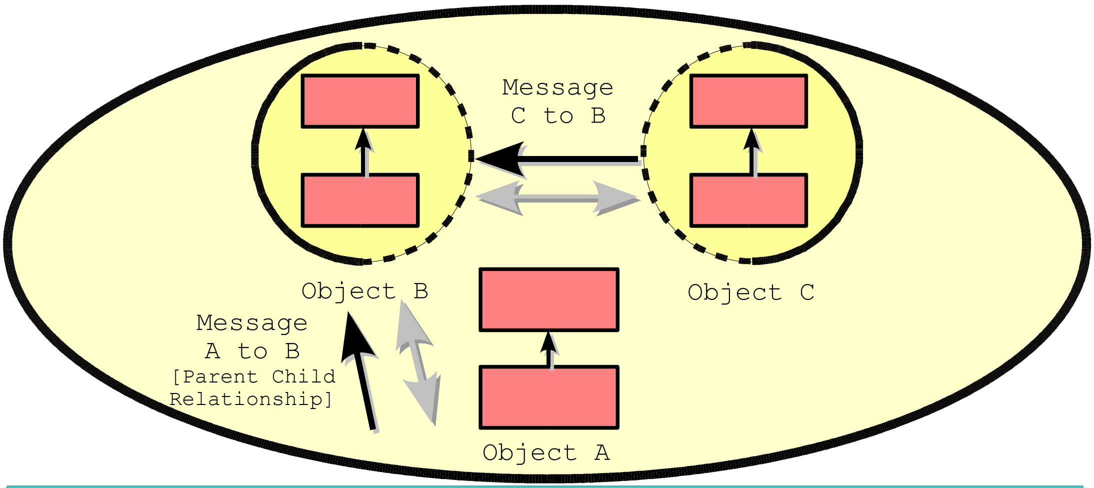

Note:
1) Solid lines have been used to show the direction of single messages.
2) Grey lines show the possible communication directions.

Figure 3.2

Classes are often referred to as the *blueprints* that define objects. They provide the schema which defines an object.
E.g. The class 'person' can be instantiated to create objects 'Adam' and 'Emma', of type 'person', which inherit the properties (*attributes and methods*) from the 'person' class.

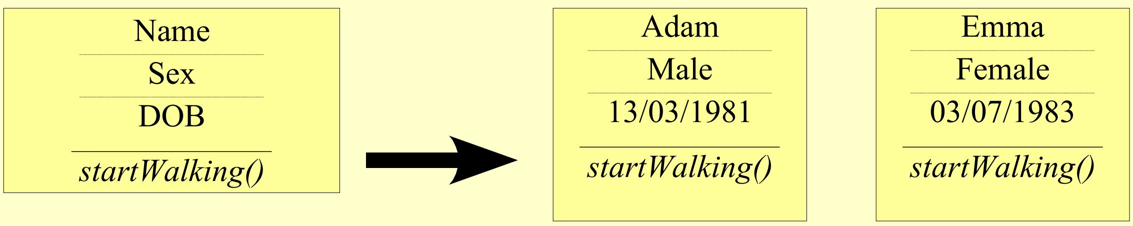

Note: The terms Objects and Classes are used interchangeably.

3.1.2 Abstraction mechanisms

The object-oriented data model consists of two abstraction mechanisms:

- ***Generalisation***

  A class is able to be a sub-class of another object class, e.g. it is possible to define Dog and Cat as sub-classes of the Animal class.
- ***Aggregation***

  A class can consist of lower-level classes. E.g. a robot may consist of 'head', 'arm', 'torso' and 'leg' classes.

3.1.3 Inheritance

There are two types of inheritance:

- ***Structural inheritance***

  A sub-class inherits the attributes of a superclass (*also known as its parent-class*). E.g. Suppose the Animal class has attributes 'Size' and 'Colour', then the sub-class Dog inherits those attributes.
- ***Behavioural inheritance***

  In a similar fashion to attributes, a sub-class is able to inherit methods (*actions*) from its superclass. E.g. The Animal class may consist of a method 'eat()'; this would be inherited by sub-classes such as Dog.

3.2 The World Wide Web as an OO System

It has long been the goal of the W3C to exploit the benefits of Object Orientation in the context of the WWW. They have developed a number of OO systems including Object-Oriented XML 2.0, Web Services Architecture*, OWL Web Ontology Language and SOAP†, primarily to aid computational systems to extract information. However, it has been difficult to locate W3C work describing the application of OO systems and categorisation techniques for altering the structure and conceptual model of the WWW. This is possibly because the W3C see that the future of the WWW lies in **computation agents** that are able to seamlessly traverse the **semantic web** to retrieve the required data. This area has long been the focus of research and applications for well understood reasons. Consequently, this research will investigate alternative approaches with an emphasis on human collaboration, computer collaboration and distributed computing.

---

* "*A Web service is a software system designed to support interoperable machine-to-machine interaction over a network. It has an interface described in a machine-processable format (specifically WSDL). Other systems interact with the Web service in a manner prescribed by its description using SOAP messages, typically conveyed using HTTP with an XML serialization in conjunction with other Web-related standards.*" W3C (2004b)

† SOAP used to be an acronym for Simple Object Access Protocol, however, according to W3C (2003b) this is no longer the case.

GRID computing* is a form of parallel computing that can encompass some of the OO concepts.

A Grid can be defined as the combined computing resources (*e.g. processing, storage, bandwidth*) that as a whole can be utilised. This can provide immense computational power as demonstrated by projects such as Oxford University's "Screensaver Lifesaver" which utilises GRID technology for cancer research, see University of Oxford (2004).

Summarised from: Globus Alliance (2004) and Pham et al. (2004, p.3)

Pham *et al.* (2004) describe "*The Integration of Peer-to-peer and the Grid to Support Scientific Collaboration*". They state that scientific communities can benefit from the collaboration of resources through both peer-to-peer and Grid technologies.
They propose an architecture of two layers:

- Computation layer: a placeholder for computational resources, i.e. Data storage, CPU Processing, database, etc.
- Collaboration layer: The collaborative environment for the front end user, using a community approach for delivering and sharing services, and other types of resource.

An example to help illustrate this concept is shown below.

E.g.

A chemist in Combustion Chemistry has found a new way to calculate a parameter for experimental data. He has implemented his method in a software program on the Grid. In the Grid environment, the software program is exposed as a service. Now he wants to share his new way of calculation with his colleagues.

Through the collaboration layer, he publishes an advertisement about his new services to other peer members. If a member is interested in the service, that member can contact the chemist for authorisation. Then the member can get access to the Grid and execute the service.

Source: Pham *et al.* (2004, p.9)

The work of Pham *et al.* (2004) links the integration of peer-to-peer and the Grid with the Semantic Grid, see Figure 3.3.

The diagram shows two perspectives; Part A illustrates different peer-to-peer communities (*collections of nodes*) which have formed different ontologies, see part B.

*"*Grids are persistent environments that enable software applications to integrate instruments, displays, computational and information resources that are managed by diverse organizations in widespread locations.*" Globus Alliance (2004)
"*a computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities. An infrastructure that enables flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions and resources.*" Pham et al. (2004, p.3)

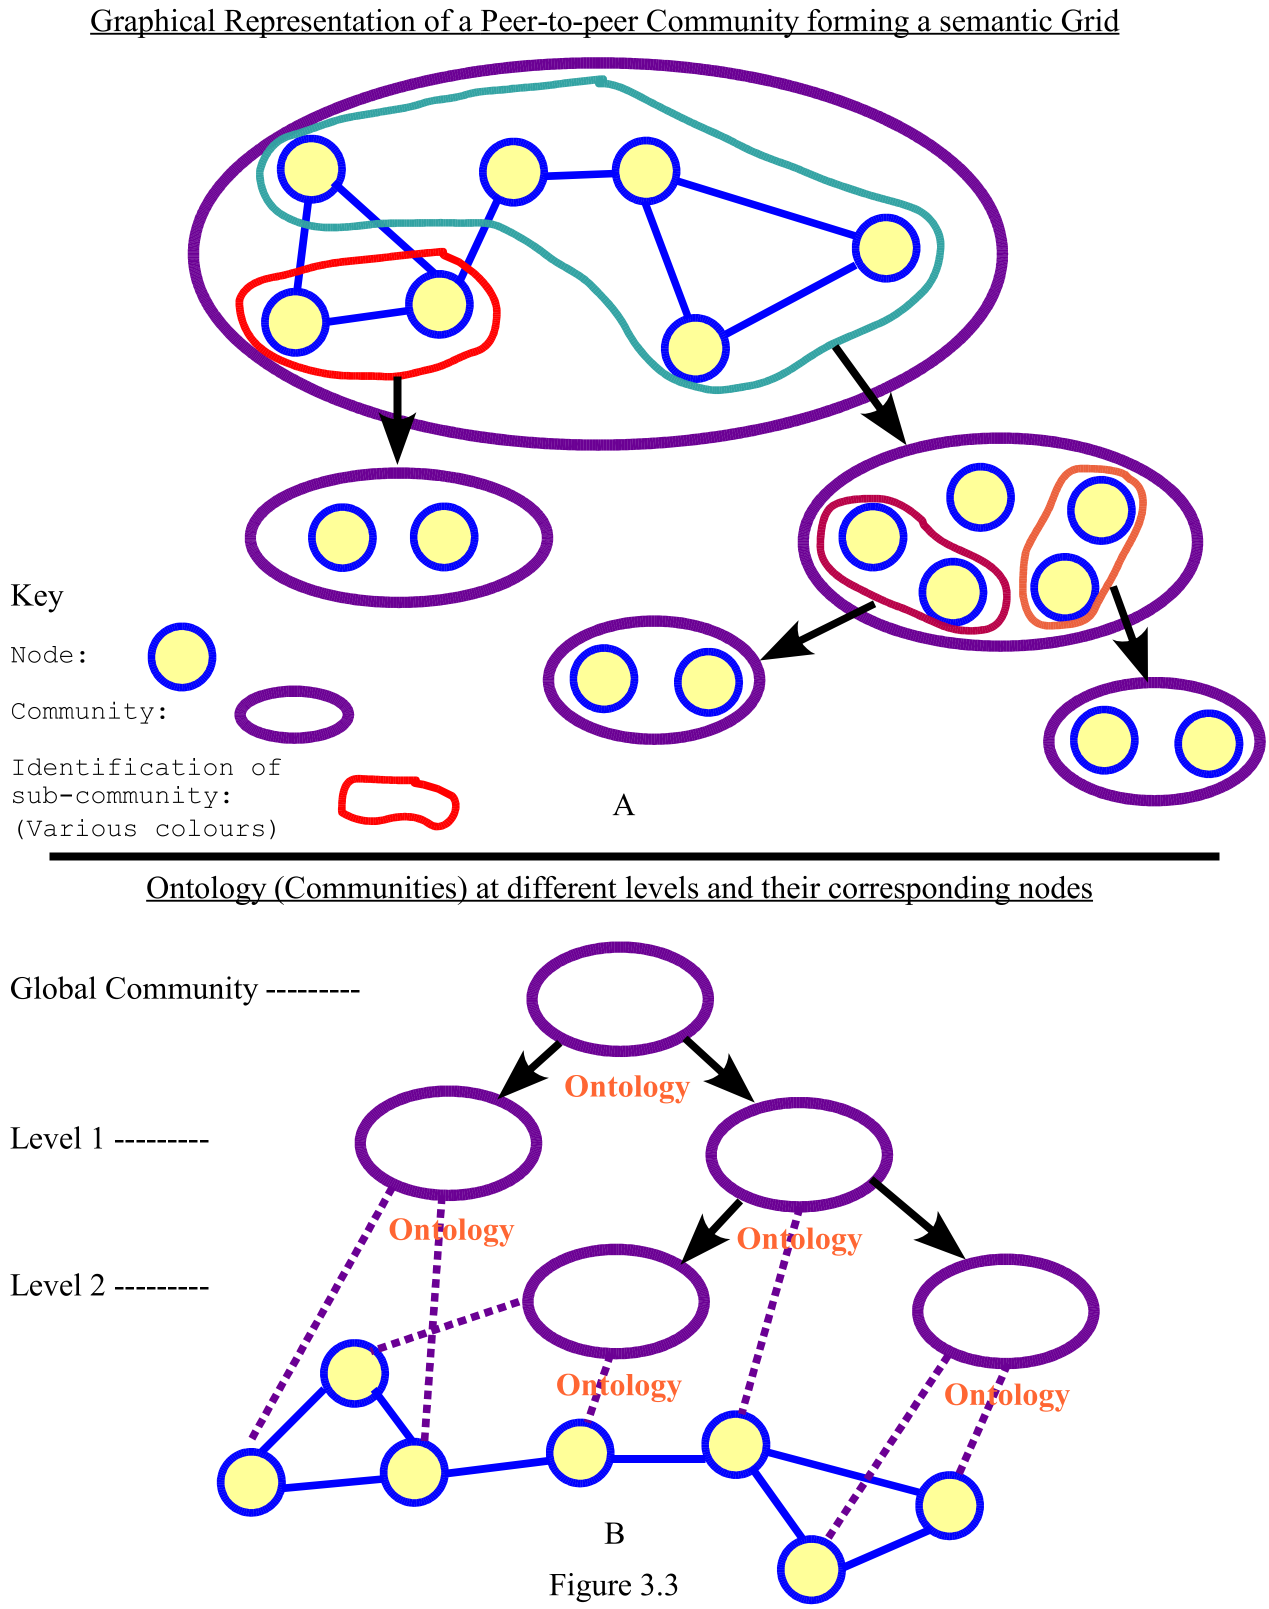

Figure 3.3

All diagrams have been adapted from: Pham *et al.* (2004, p.11 - 14)

One criticism of Pham *et al.* (2004) is that they don't illustrate their concept particularly well in terms of OO. Fortunately their work can be extended relatively easily to accommodate OO, see the basic example in Figure 3.4, illustrating the "Book Repository" super object, the category objects and the individual book nodes. Their availability can be likened to e-Books available on the WWW.

Graphical Representation of a hierarchical, peer-to-peer system to form objects

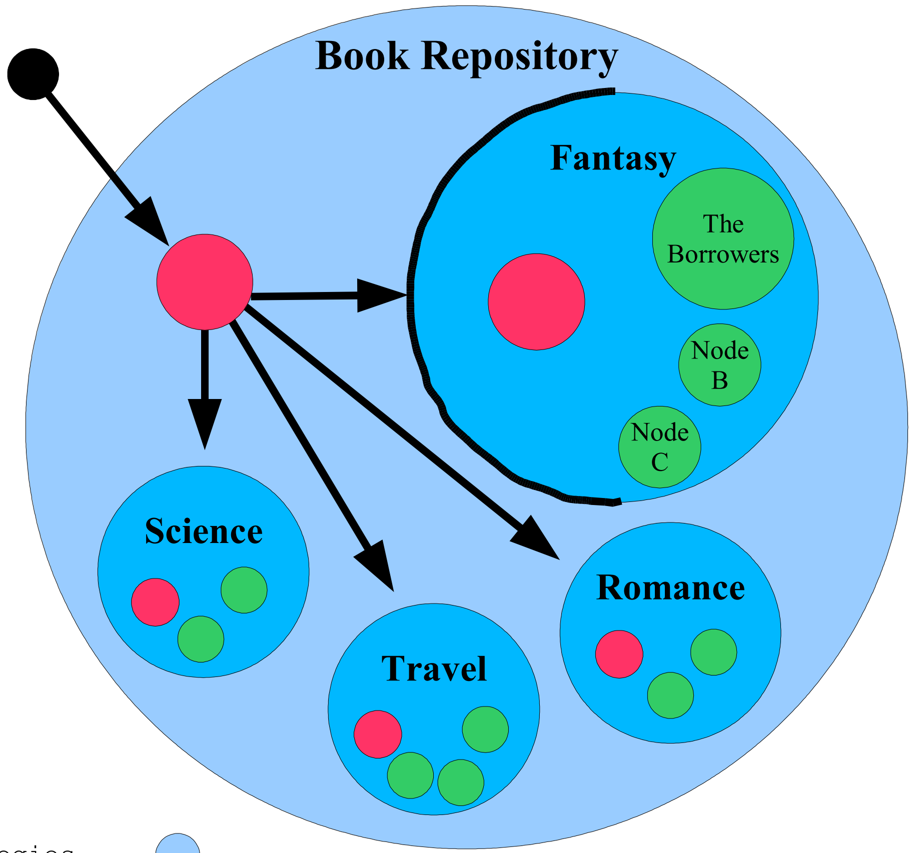

Key

```
Objects/Groups/Ontologies

Node Lookup Services (NLSs)      <-- Implementation
                                     It is more likely that this service would
                                     run on multiple nodes.
                                     This is discussed in later sections.
NLS Root Domain Server

Nodes

Interface (only shown on the fantasy object):

Note:
* The ontologies (Objects) have attributes, methods and interface codecs.
* The objects could theoretically be public, private or protected.
  Furthermore, they (and the nodes) can be seen as products or services.
* A NLS is aware of the address for each internal object's NLS
  (similar to the current Domain Name System),
* The object is supported by the combined computational power each node within
  provides (CPU cycles, storage space etc.) i.e. they form a Grid.
* The NLS' could theoretically provide a range of facilities from mere node
  IP/address resolution to complex search facilities.
```

Figure 3.4

Further explanation is required to show how this example is related to the work provided by Pham *et al.* (2004). With reference to the WWW, the "community" that Pham *et al.* describe, are the servers that provide the web-pages (*web-servers*). It is these that provide the computational resources required to support the objects. Together they form a Grid creating an ontology, as shown in Figure 3.3. The nodes within the object are related by the subject that they are concerned with. It is evident that the objects (*in this methodology*) could easily incorporate the abstraction and inheritance aspects associated with OO, e.g. it is quite possible for "The Borrowers" node to inherit certain properties from the book object such as formatting (*e.g. html*), read-only access etc.

Bibliography

Beynon-Davies P. (2004). *Database Systems* Third edition, Basingstoke: Palgrave.

Chuanfeng, C. and Zengzhi, L. (2000) Internet Object-oriented Information Model, *Technology of Object-Oriented Languages and Systems*, International Conference, 285-290.

Connolly, D. and Soley, R. M. (1996) *Joint W3C/OMG Workshop on Distributed Objects and Mobile Code,* Massachusetts, Massachusetts Institute of Technology.
Available from: http://www.w3.org/OOP/9606_Workshop/, [Date Accessed: 26 August 2004]

E-Science (2004) *GGF Semantic Grid Research Group (SEM-RG)*, Southampton, University of Southampton.
Available from: http://www.semanticgrid.org/GGF/ggf11/, [Date Accessed: 17 July 2004]

Francis, P. (2003) Is the Internet Going NUTSS?, *IEEE Internet Computing*, 7 (6), 94-96.

Globus Alliance (2004) *The Globus Alliance*, California, The Globus Project.
Available from: http://www.globus.org, [Date Accessed: 17 July 2004]

Nilsson, J. (2002) A Pure Object-oriented Domain Model by a DB Guy, Part 1, Connecticut, Jupitermedia Corporation.
Available from: http://www.devx.com/vb2themax/Article/19892/0/page/1, [Date Accessed: 15 July 2004]

Pham, T. V, Lau, L. M., Dew, P. M. (2004) *The Integration of Peer-to-peer and the Grid to Support Scientific Collaboration*, Leeds, University of Leeds.
Available from:http://www.semanticgrid.org/GGF/ggf11/GGF11-final.ppt, [Date Accessed: 17 July 2004]

Sun Microsystems (2004a) *What Is a Message?*, California, Sun Microsystems Inc.
Available from: http://java.sun.com/docs/books/tutorial/java/concepts/message.html, [Date Accessed: 15 July 2004]

Sun Microsystems (2004b) *Grid Adoption Trends*, California, Sun Microsystems Inc.
Available from: http://www.sun.com/aboutsun/media/presskits/sgeee/gat.gif, [Date Accessed: 16 July 2004]

Thomson Learning(2004) *O Glossary page*, Kentucky, Thomson Corporation.
Available from: http://www.scsite.com/dc/main_glossary_o.htm, [Date Accessed: 15 July 2004]

University of Oxford (2004) *Screensaver Lifesaver - University of Oxford*, Oxford, University of Oxford.
Available from: http://www.chem.ox.ac.uk/curecancer.html, [Date Accessed: 17 July 2004]

W3C (1999a) *WWW and OO: Objects on the web*, Massachusetts, Massachusetts Institute of Technology.
Available from: http://www.w3.org/Consortium/Prospectus/OO.html [Date Accessed: 26 August 2004]

W3C (1999b) *Schema for Object-Oriented XML 2.0*, Massachusetts, Massachusetts Institute of Technology.
Available from: http://www.w3.org/TR/NOTE-SOX/, [Date Accessed: 26 August 2004]

W3C (2003a) *Web Services Architecture*, Massachusetts, Massachusetts Institute of Technology.
Available from: http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/, [Date Accessed: 26 August 2004]

W3C (2003b) *SOAP Version 1.2 Part 1: Messaging Framework*, Massachusetts, Massachusetts Institute of Technology.
Available from: http://www.w3.org/TR/2003/REC-soap12-part1-20030624/, [Date Accessed: 26 August 2004]

W3C (2004a) *RDF Primer*, Massachusetts, Massachusetts Institute of Technology.
Available from: http://www.w3.org/TR/REC-rdf-syntax/, [Date Accessed: 17 July 2004]

W3C (2004b) *Web Services Glossary*, Massachusetts, Massachusetts Institute of Technology. Available from: http://www.w3.org/TR/2004/NOTE-ws-gloss-20040211/, [Date Accessed: 26 August 2004]

W3C (2004c) *OWL Web Ontology Language Overview*, Massachusetts, Massachusetts Institute of Technology. Available from: http://www.w3.org/TR/owl-features/, [Date Accessed: 26 August 2004]

Section 4: Investigation of the Realm concept with respect to the conceptual structure of the Internet

4.1 Realm Introduction

Section 3 provided a grounding and introduction to OO and the WWW. This concept, however, is not limited to the WWW as it can be applied to many aspects of the Internet and other telecommunications systems. The application of OO in this way is known throughout this documentation as **Realms**, and each group is known as a **realm**. We provide the following definition:

> *A realm may be defined as the interconnection (cluster) of devices to form an ontology based on semantic reasoning. The realm concept is an abstraction method to allow the identification of nodes logically rather than arbitrarily*.

For example, a fishing web page can be located in a public fishing realm, where all fishing related material can be located. A node in this context is NOT necessarily a computer, but can be a computational resource available at a specific location (*e.g. Web-page or streaming video*). The purpose of the realm concept is to enhance the virtual structure and NOT necessarily physical performance, although the latter may be a by-product in some instances. Further to this, a realm specifies the communication methods and supports inheritance and generalisation mechanisms. This could be used to promote the existence of "loosely coupled"* systems.

Two important properties of the Realm are: logical existence and standardised interface.

It is important to make a distinction between information and interface web-pages. The former merely presents information in a similar way to paper whereas the latter provides an interface to access an abstract system, e.g. a Short Message Service (SMS) sending web-page is an interface web-page as it provides an interface to an SMS sending system.

It is evident from the W3C publications that their work covers a diverse range of topics. It is possible that the W3C will develop an information structure that extends the WWW's scope, suppressing any of the "realm" benefits.

4.2 Realms and the Structure of the Internet

The Realms concept could dramatically enhance the conceptual model that users have of the Internet. It could provide an interface that enable users to seamlessly locate nodes based on meaning, and interact with nodes in the most appropriate way. Realms need not replace any of the current systems, but rather act as an enhancement, see Figure 4.1.

This means that (*initially at least*), the Internet structure would NOT be altered, i.e. once a node has been located, from within a Realm, the Domain Name/IP/URL is established and contact can be made in the standard way. However, in the future, it may be possible, using Grid technology

* Loose coupling is "*The friction-free linking enabled by web services (or any Service Oriented Architecture). Loosely coupled services, even if they use incompatible system technologies, can be joined together on demand to create composite services, or disassembled just as easily into their functional components. Participants must establish a shared semantic framework to ensure messages retain a consistent meaning across participating services*". LooselyCoupled.com (2004)

(*parallel computing*), to provide the node "virtually" from the Grid resources. One example of this type of technology is Freenet (2004), a project that provides a totally decentralised network and enables entire websites (*amongst other data*) to exist using decentralised computers (*typically PCs*). Freenode relies on standard Internet users to meet the bandwidth, storage and processing requirements.

Graphical Representation of how the Realms concept fits into the the Internet's abstraction model

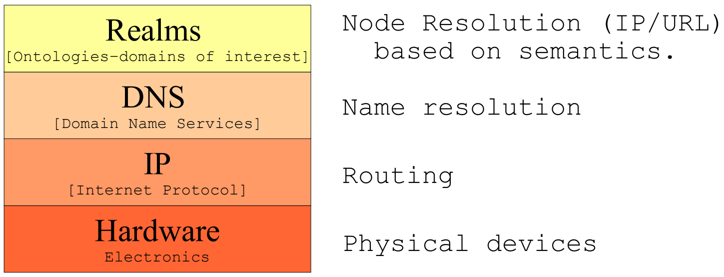

Figure 4.1

The Realms concept as described so far can be understood from studying Figure 4.2. A key point is that the node need not be a physical device, but rather be a location on a device.

Diagram illustrating a collection of nodes and their relationship to the different realms

- a **single node** can exist in **multiple realms**,
- a **node** is part of a computing resource (e.g. Server) BUT may be a **location** rather than the device itself (e.g. URL).

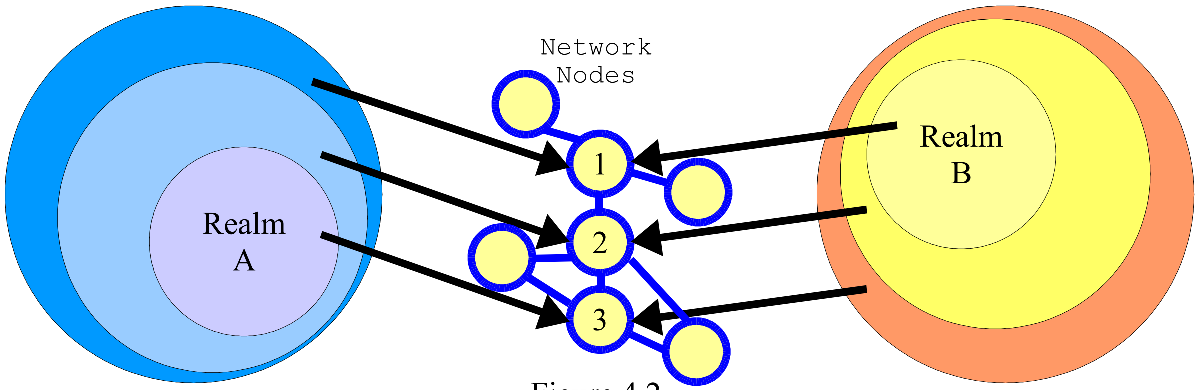

Figure 4.2

What happens to the realm should a node go down?
The node is obviously part of the Grid that supports a realm. Could other computers detect this and essentially drop information relating to the node? Li *et al*. (2002) identify, with reference to peer-to-peer (P2P) networks, one method to solve this problem. When a node unexpectedly exits from a P2P network, the failure of the node must be detected by the directly connected nodes. Further to this, the routing tables that use the node as a soft router, to relay packets, need to be rebuilt to accommodate the change.
However, if the node was unexpectedly lost (*e.g. it crashed or there was a network problem*), would it be possible for the Grid to temporarily support the disconnected node?

4.3 The Internet, Realms and the potential scope of devices

It is evident, from sections one and two, that the scope of devices that are connected to the Internet is likely to increase dramatically and that there are two major strands to this: the physical devices (*Personal Computers, fridges, sensors etc.*) and virtual (*intangible*) devices (*music songs, films, software programs etc.*)

4.3.1 Physical Devices

This may be more useful for business users than domestic users, e.g. to locate/interact with pipeline sensors.

As more physical devices connect to the Internet, it is likely to become more difficult to locate a node without knowing its unique identifier. The realm solution is to create a virtual group of like devices to form ontologies, based on meaning. This concept could be applied to all manner of physical devices, although it is assumed that only devices that would benefit from, or be beneficial to, the Internet as a global information system would be included, i.e. it would be pointless to create ontologies consisting of physical devices that would provide little or no benefit.

One example of a type of physical device that could benefit from this concept is telephones. It would be possible to systematically group them (*assume IP phones for simplicity*) according to their geographical details shown in the public phone book. This would enable efficient interrogation by any device understanding the realm's protocol, by limiting the domain of interest to a specific realm such as a town or city (*this could be exploited to provide incredible data mining capabilities by using devices employing this concept*). Figure 4.3 provides a graphical representation for the example; Figure 4.4 provides an indicative context for the realms (*objects*) employed in the example.

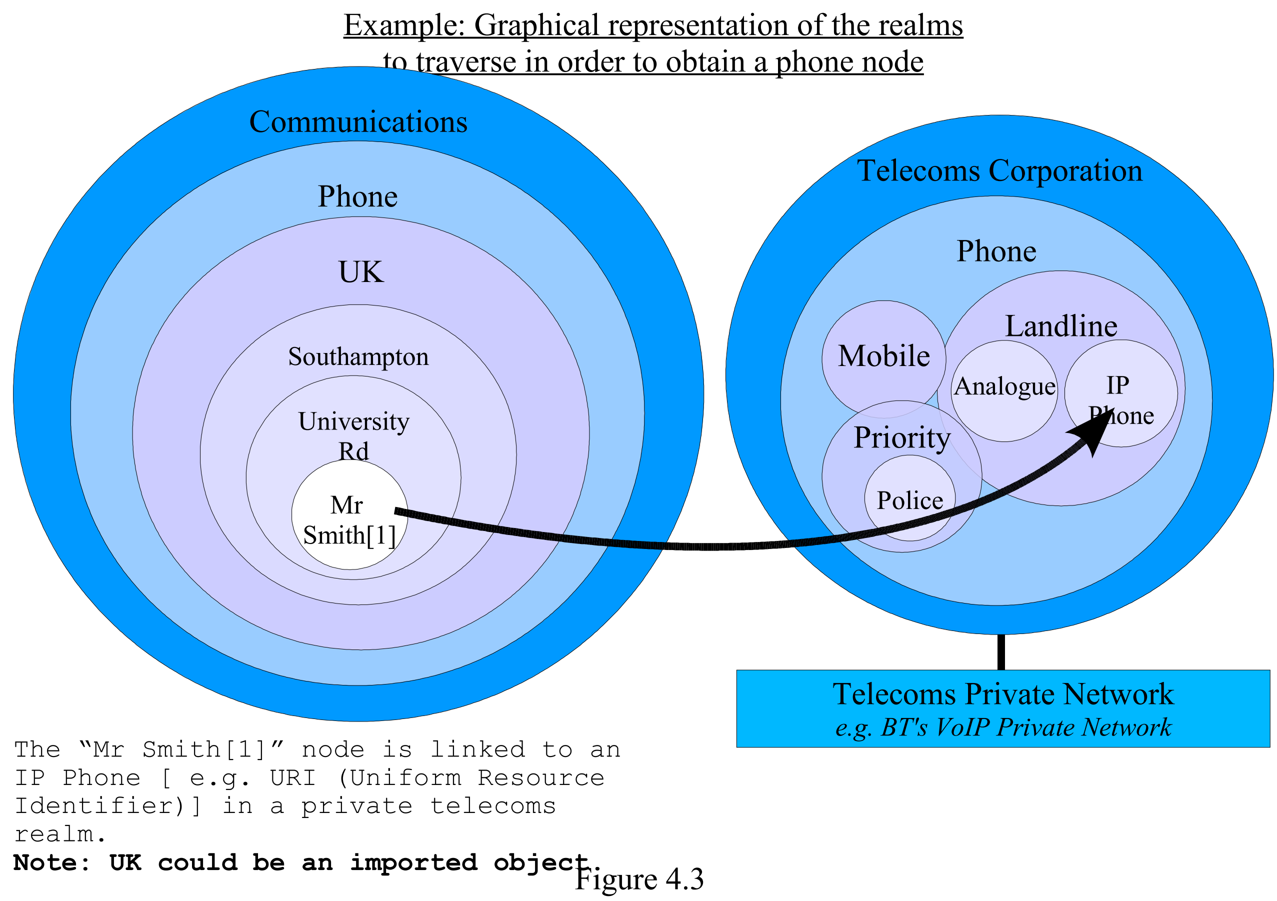

Figure 4.3

It must be reiterated that a node, such as the IP Phone in Figure 4.3, can exist in multiple locations. Thus the Mr Smith[1] node in Figure 4.3, could be in any number of realms provided it has the computational resources to support itself (*e.g. it could be listed in the public South African and private Smith Family realms*).

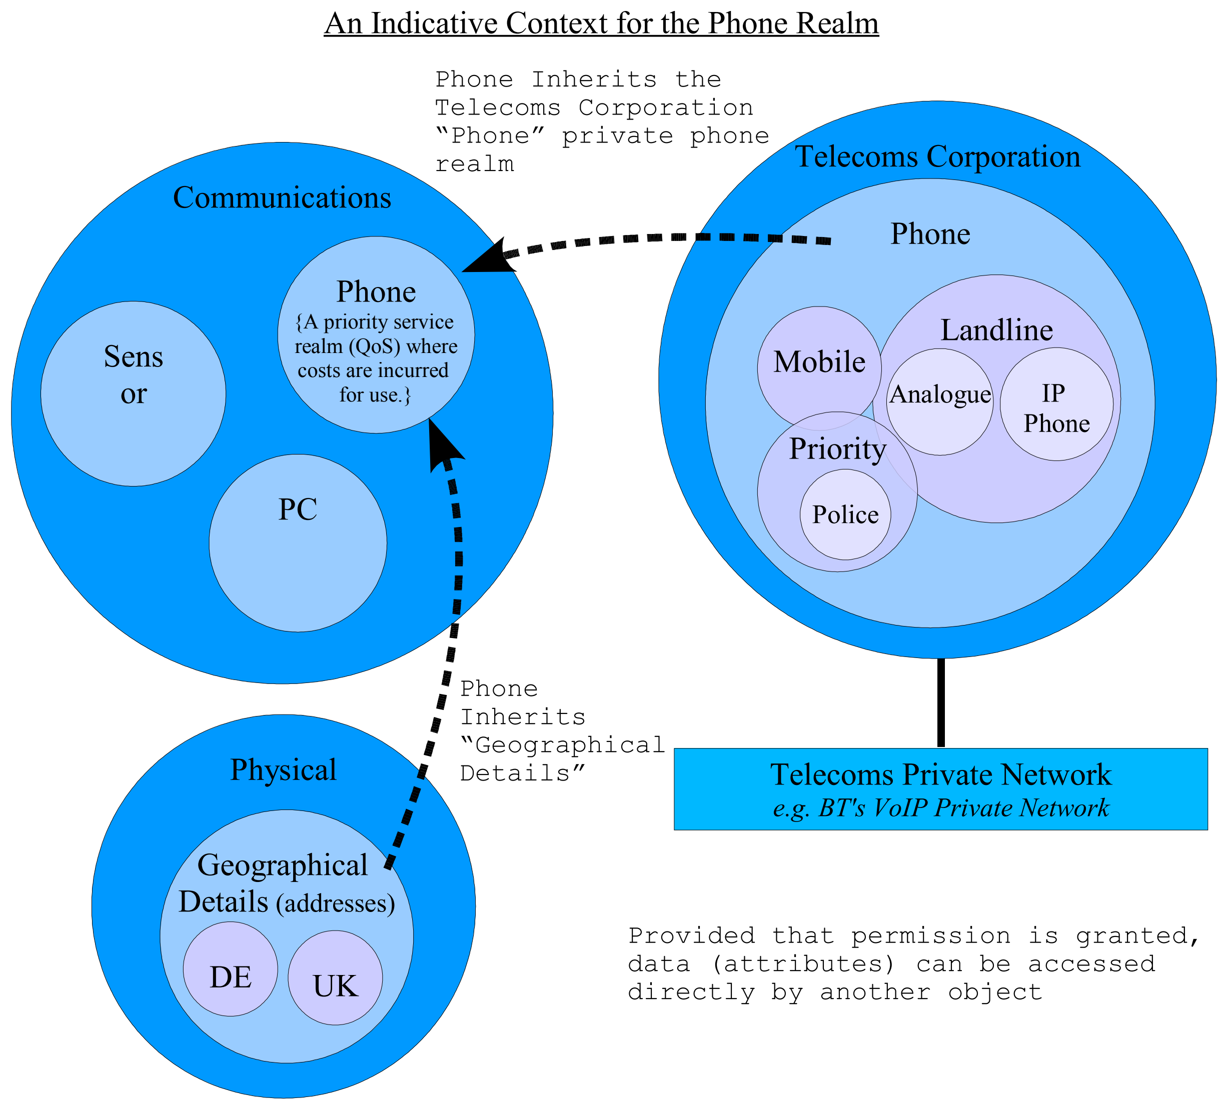

Figure 4.4

4.3.2 Virtual Devices

As defined previously, a virtual device is a resource/service that is provided by a computerised system such as a server. This can include anything that provides an interface to a system and can range from an artificially intelligent "computer doctor", where questions are answered and recommendations given, to electronic books and films, where they are simply accessible for streaming/download.

However, despite the physical difference, devices under this category could be "organised" in a similar way to physical devices, based on logical reasoning. For example, the film "Love Definitely" may be organised as shown in Figure 4.5, in addition to any number of other realms.

Example: Graphical representation of the realms
to traverse in order to obtain a film node

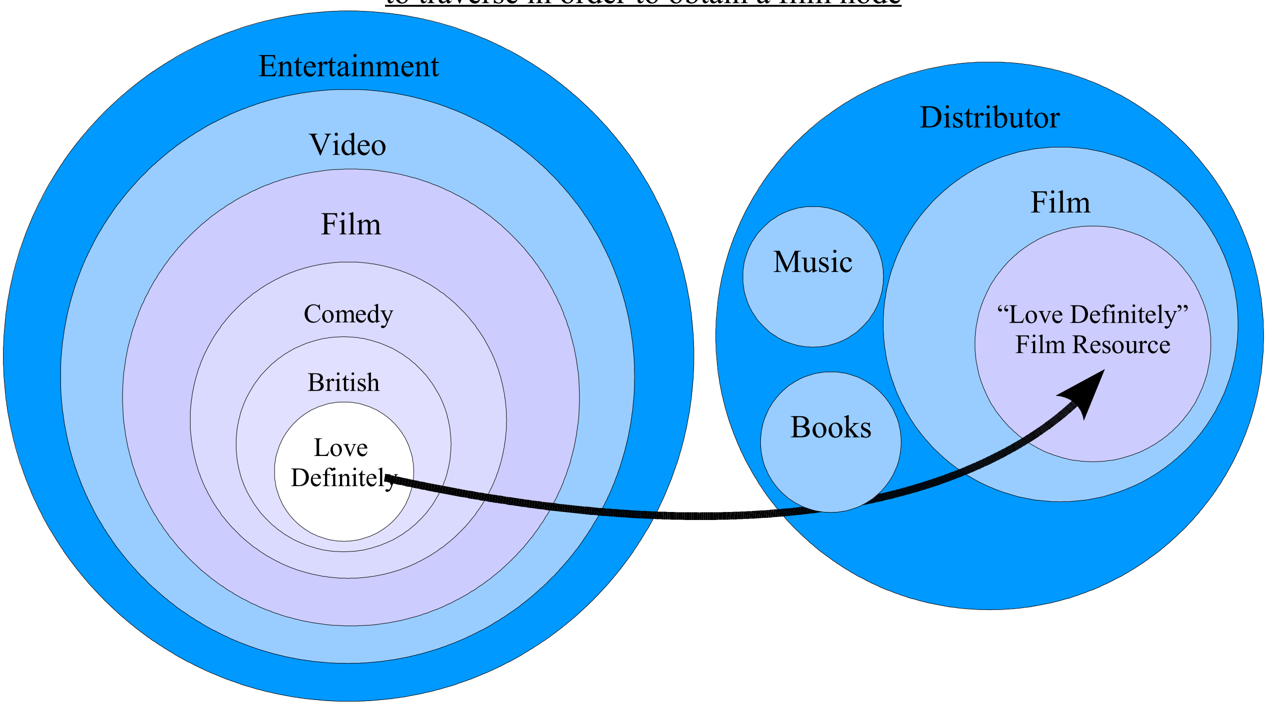

Figure 4.5

4.3.3 Quality of Service (QoS)

This discusses the problem with reference to phones, but applicable to any situation where QoS is important.

Although IP telephony was discussed above, the Realms concept was only discussed with reference to phone location. Another potential property of the Realms concept is that QoS could be ensured. Phone companies could combine their resources (*bandwidth/telecoms lines*) to create one real-time bandwidth that any system, accessing the realm, could make use of in return for a fee. Assuming QoS within the realm is perfect, the only aspect where QoS could be affected is access points in and out of the Realm. In the case where communication does not cross the interface borders, there is no loss in QoS, e.g. if a phone device is connected directly to the realm in a similar way to the traditional land-line phone.

A final point is that it would be possible for devices to access the real-time "phone realm" ad-hoc; e.g. whenever phone services are required, software contacts the ISP informing them of the request, and the ISP reserves bandwidth and ensures the customer has direct assess to the "Phone" realm. Further to this, a similar system could be employed for streaming video, where the customer requests a resource "on demand" and is given real-time resources to ensure QoS (*appropriate jitter and bandwidth*) is maintained in return for a fee. See Figure 4.6 for a graphical representation of this example.

This feature is not limited to the Realms concept as independent phone companies could provide servers and leased lines to help ensure QoS, but this is likely to be company specific. With "Realms", real-time bandwidth could be accessible by any computer node or software system that has the relevant credentials (*e.g. has paid a fee*) and isn't limited to a prior agreement in the same way as dedicated leased lines. Special realm nodes (*or a realm service*) could be used to record usage and ensure fair distribution between the telecommunication providers.

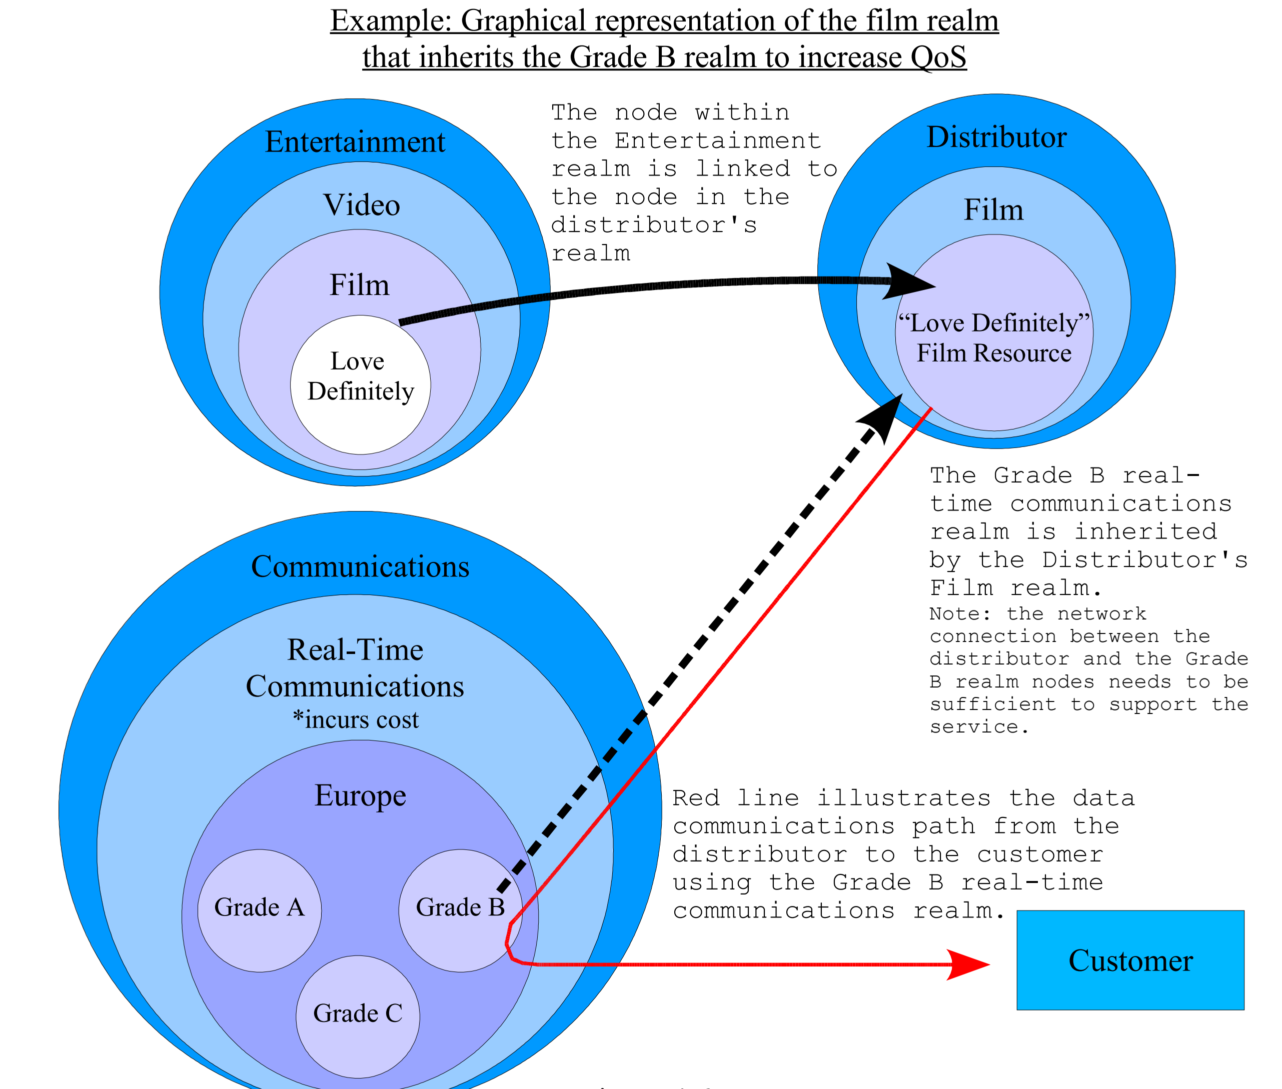

Figure 4.6

4.3.4 Security

Realms could provide added security in three distinct ways:

1. By authenticating clients (*devices or users*) to use the realm.
2. By providing authentication of the devices to become peers within a realm.
3. By allowing communication companies to combine communication resources to create secure communication lines in a similar way to the Grade B communication realm in Figure 4.6, e.g. it would be possible to have a bank realm, where data is only transferred over secure communication lines between source and destination.

4.3.5 Other Contexts

The realm concept could help widen the current scope of the Internet by providing added value. It could help put Internet resources into context, and help standardise interfaces between similar nodes. This means that similar objects could exist, and remain useful because there is some virtual “common ground” which binds the nodes, and the standardised interface could allow other systems to exploit the objects using the common realm protocol.

For example, the UK county councils are essentially different instances of the same organisation and yet their web-site information is typically obtained from different locations. A realm could help disguise the differences because all nodes within the realm would conform to the same standards (*e.g. Methods*). This could give rise to more abstract systems. For example, Towns/Councils within the UK could, collaboratively, create a UK town realm which might specify that nodes provide emergency contact information, the town’s history, a town map, accommodation details, special events etc. (*tourist information*). The common interface could enable system developers to exploit the realm easily, e.g. car manufacturers could easily create systems that make use of the realm facilities to display information within the car because of the standardised interface, see Figure 4.7.

Example: Graphical representation of the UK Town realm and the abstraction (interface) mechanism

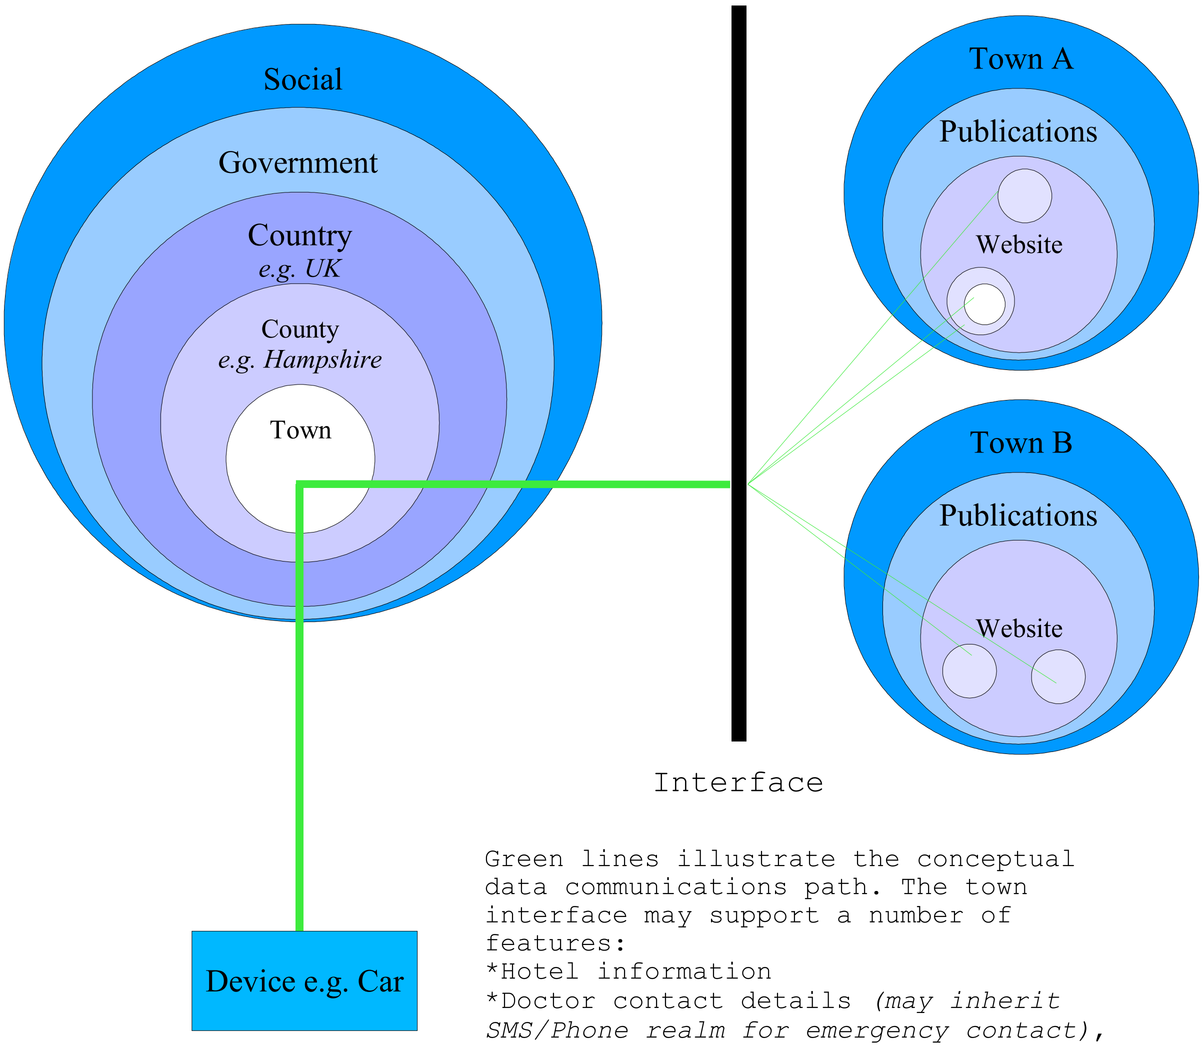

Figure 4.7

It is important to state that we do not suggest that the example in Figure 4.7 would be a viable system, but are merely using it to describe the principle of abstraction.

4.3.6 Drag and Drop Resources

Realms could help alter the focus from the installation of programs and the procurement of devices, to the installation and procurement of resources. It could be possible to "drag and drop" realm interfaces onto a device, thus increasing the device functionality (*e.g. map access*). Therefore, if one has a PDA and wants to provide added functionality (*such as phone facilities*), it could be possible to "drag and drop" the Phone realms interface (*which could be tailored to the device capabilities*) onto the device.

Bibliography

Freenet (2004) *The Freenet Project - news – beginner*, California, Open Source Development Network Inc.
Available from: http://freenet.sourceforge.net/, [Date Accessed: 19 July 2004]

Li, Y., Li, J., Yu, K., Wang, K., Li, S., Zhang, Y. (2002) A Peer-to-peer Communication System, China, Microsoft Research Asia and Xi'an Jiaotong University.
Available from: http://research.microsoft.com/asia/dload_files/group/imedia/2002p/pcm02-jiangli.pdf, [Date Accessed: 26 July 2004]

LooselyCoupled.com (2004) *Loosely Coupled glossary*, Somerset, Procullux Media Ltd.
Available from: http://looselycoupled.com/glossary/, [Date Accessed: 28 July 2004]

Mahalingam, K. and Huhns, M. N. (1997) A Tool for Organizing Web Information, *IEEE Computing*, 30 (6), 80-83.

W3C (2000) *WWW and OO: Objects on the web*, Boston, World Wide Web Consortium.
Available from: http://www.w3.org/Consortium/Prospectus/OO.html, [Date Accessed: 25 July 2004]

W3C (2004a) *RDF Semantics*, Boston, World Wide Web Consortium.
Available from: http://www.w3.org/TR/2004/REC-rdf-mt-20040210/#intro, [Date Accessed: 25 July 2004]

W3C (2004b) *OWL Web Ontology Language Overview*, Boston, World Wide Web Consortium.
Available from: http://www.w3.org/TR/2004/REC-owl-features-20040210/, [Date Accessed: 25 July 2004]

Section 5: Interviews with potential end-users and experts

5.1 Methodology

It was important to obtain opinions from potential end-users because they often view systems in different ways and can provide valuable feedback. Five computer users of different abilities were approached and later interviewed for this research.

A semi-structured interviewing technique was employed in order to extract as much relevant information within the limited time and resource constraints. Preece *et al.* (2002, p.211) states that interviews are "*good at getting people to explore issues...interacting with a human rather than a sterile, impersonal piece of paper or electronic questionnaire encourages people to respond, and can make the exercise more pleasurable*".

The interviews took the following form:

1) Explanation – The author explained to the end-users the major aspects of all prior sections. This focused on the diversity of devices, the likely uses of Internet technology in the future, and the benefits of object orientation.

2) Questionnaire – The author asked a number of questions. Each participant was given the opportunity to comment and request further explanation throughout their interview.

The interview with the two computing experts was performed together and was deliberately less structured than with the potential end-users. This provided a flexible approach and enabled the experts to discuss the aspects which they considered to be most important.

Part of the interview was used to discuss other approaches to Realms.

5.2 Findings of the results of the user interviews

The following charts provide a summary of the majority of questions shown in the appendix to this section.

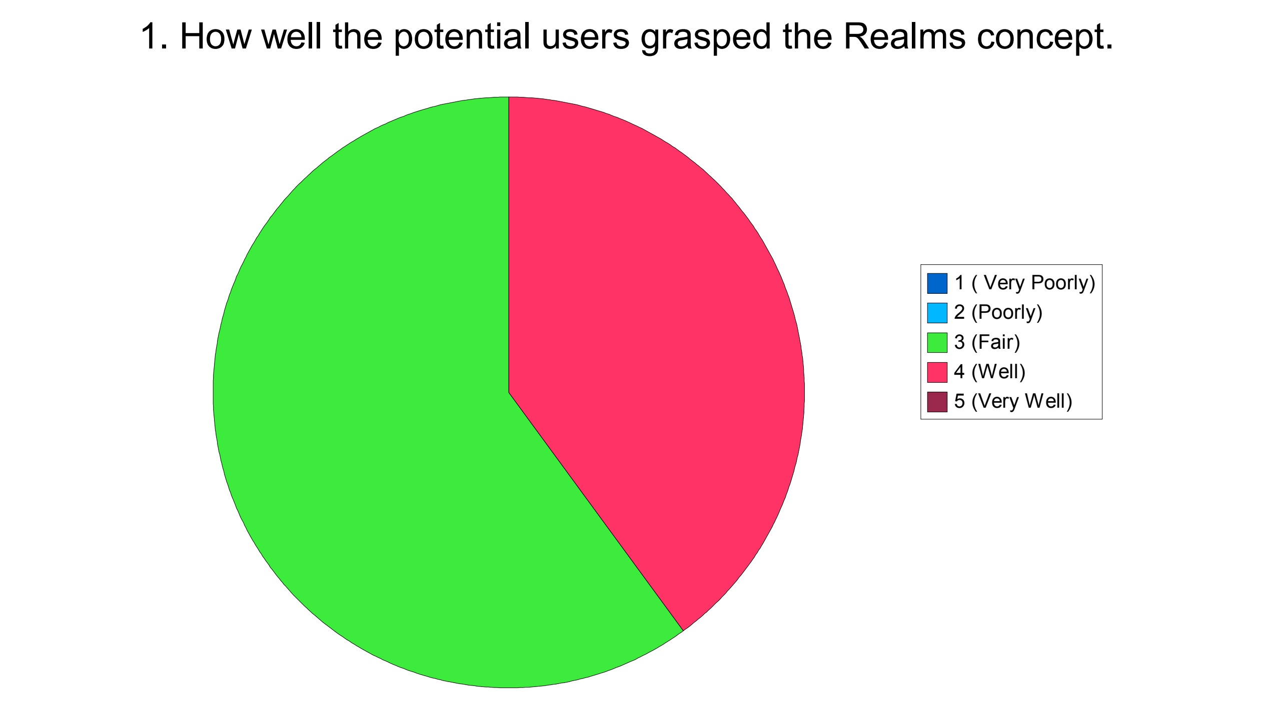

Figure 5.1

It is evident from Figure 5.1 that all the users questioned felt, after a verbal explanation, that they had a reasonable understanding of the Realms concept. This is significant because good user understanding would be crucial to the success of the Realms concept.

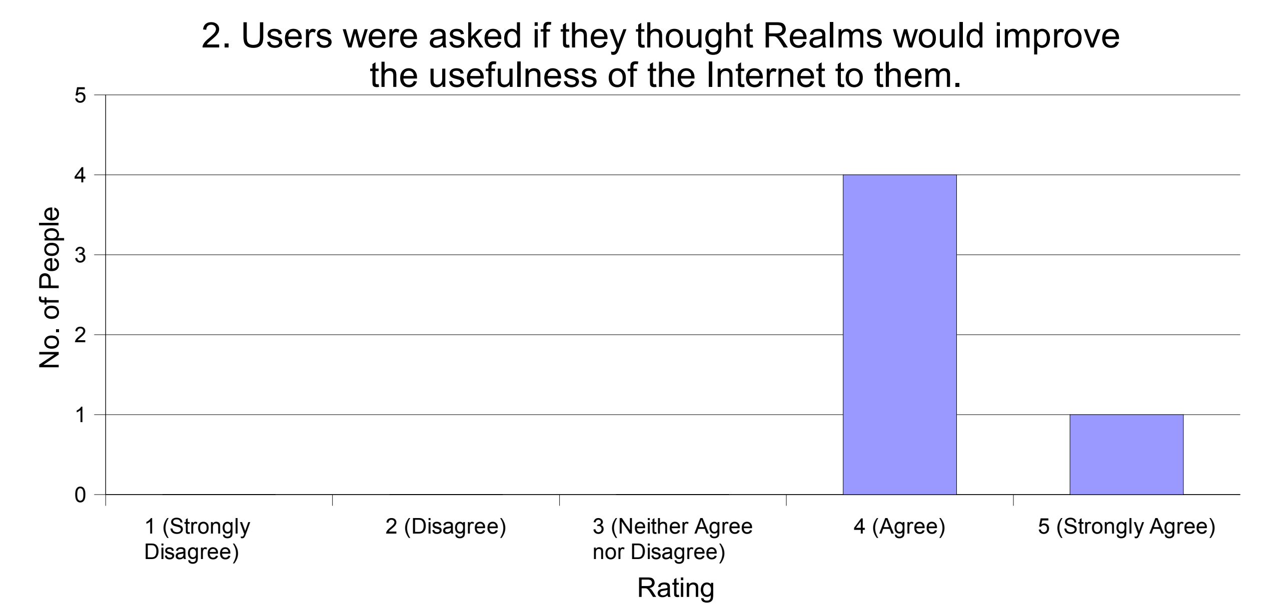

Figure 5.2

Further to merely understanding the realms concept, it is evident that those questioned thought that realms would be beneficial to them.

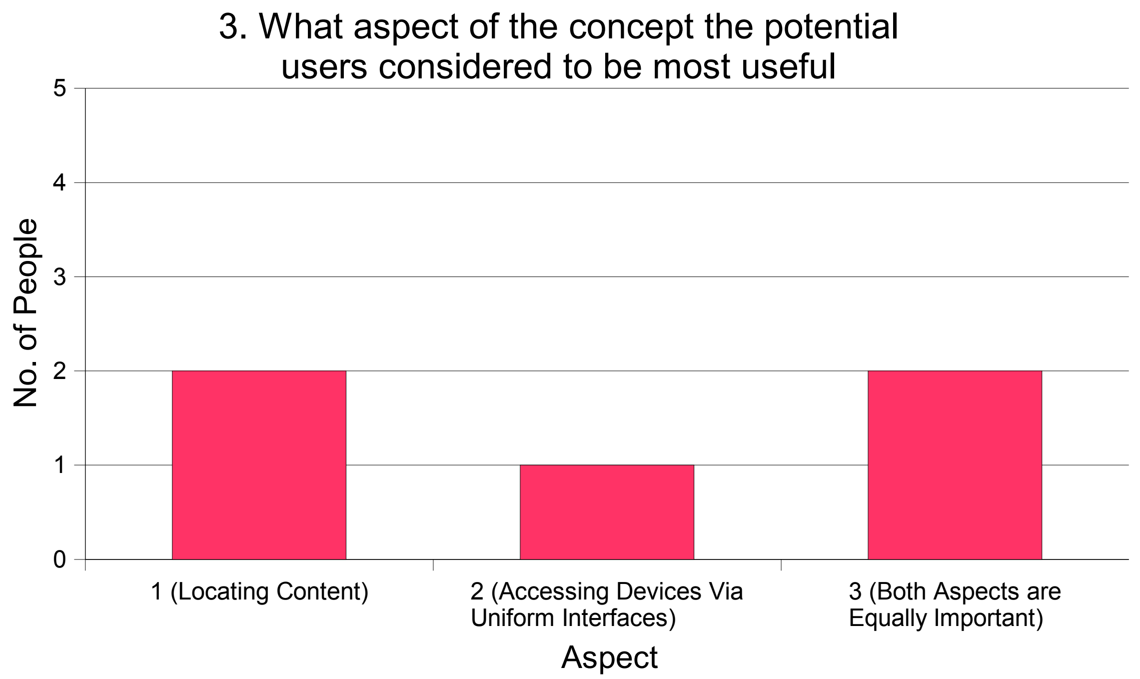

Figure 5.3

Figure 5.3 clearly shows a diverse response to this question. This is possibly explained by the different abilities of the participants and their demand for computational functionality.

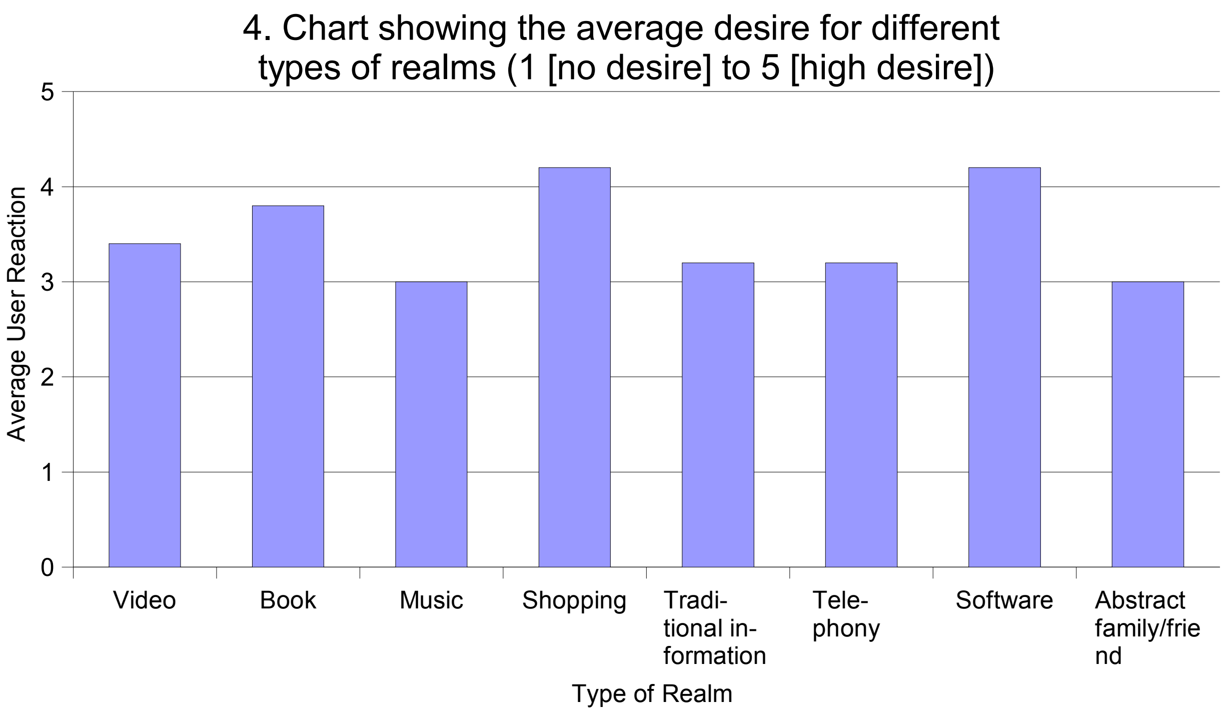

Figure 5.4

It is clear from Figure 5.4 that, collectively, the participants expressed desire for all the realms which were suggested.

The charts show that the potential end-users both grasped the concept and thought that it could provide some benefit to them. There is, however, more detailed information available from the raw results, shown in the appendix to this section. For example, participants were asked if they could suggest additional realms (*with respect to those shown in Figure 5.4*). The results in this case were quite surprising and demonstrated that those questioned applied the concept to diverse contexts.

5.3 Interviews with experts

Despite the mainly positive reaction from the potential end users of mixed abilities, two networking experts from the University of Southampton, Thompson and Weal (2004) essentially dismissed the entire concept.

Their first criticism was: “how could one be sure that this single broad concept was the most appropriate for all the different aspects of the Internet.” Further to this, they stated that past attempts to organise the Internet had largely failed, and asked why should such a concept as this be any different from those; they cited synonyms and context as likely problem areas.

Secondly, they stated that the examples of realms given (*film, music and book repositories*) didn’t provide anything extra in comparison to current systems.

They finally stated that many talented individuals are working on the various different aspects which the Realm concept provides and insinuated that one would be very naive in thinking that one concept could tackle the search, location and interface aspects. They said that there are many other ways of tackling these aspects, especially in terms of search. They stated that Realms required users to think about the context of a node within the vast Realms ontology, whereas, in their experience, users didn’t like to “think” and so the focus should be on how computers can seamlessly provide the required information.

When finally asked if either could ever envisage this concept as being important in terms of node location and interface control, they both stated that they couldn’t and that perhaps one should focus on a particular aspect rather than attempting anything of this complexity. They compared a thorough investigation of all the aspects which Realms covers in terms of research to the equivalent of four PhD theses.

Appendix

Five individuals, Hirst (2004), Howarth (2004), Maule (2004), Smith (2004), West (2004), were asked a number of questions.

1) How well do you grasp the Realms concept? {*Poorly 1 – 5 Very Well*}	
Hirst	3
West	4
Smith	3
Howarth	4
Maule	3
2) Do you think it would improve the usefulness of the Internet to you? {*Strongly Disagree* 1 – 5 *Strongly Agree*}	
Hirst	4 – Liked the idea of limiting the search based on category. He liked the fact that moderators in small realms could help reduce inappropriate nodes. In larger realms, the suggestion of voting to remove inappropriate nodes (*content*) was made (e.g. users accessing a node then realising it wasn't appropriate could "recommend" removal. Those with high removal ratings could be lowered in search ranks etc. until total removal)
West	5 – "Searching would be significantly improved".
Smith	4 – "Searching would be improved above the current methods."
Howarth	4 – "Searching would be improved and it would give an inexperienced user a foundation to locate things, e.g. the user isn't required to know specific URLS etc.".
Maule	4 – Liked the idea to limit search ("search engines provide a lot of useless hits"). Also liked the idea of realm access control and potential "secure realms" - "not necessarily for home use, but at work...".
3) If the Realms concept was beneficial, where would this concept be most useful to you? (1) Locating content (*Finding such and such*), (2) accessing devices via uniform interface, or (3) both aspects are equally important).	
Hirst	2 – Thought the uniform interface would be most useful. He felt this technology could enable him to have one global filestore that his phone, PC, MP3 player could use and for him to simply download the "storage realm" interface plug-in to each device.
West	1 – Said that the realms elimination of irrelevant nodes would be a significant improvement over the current searching methods, e.g. www.google.com.
Smith	1 – Would use it to search for data, although isn't sure whether she would make use of virtual devices.
Howarth	3 – "Would be equally useful in terms of node location and standardised interface".

Maule	3 – Said that it would be equally useful for locating information and accessing virtual devices, especially virtual devices that could provide access to experimental data.
4) From the following list of Realms, please indicate your desire for each. {*No desire* 1-5 *Highly Desire*}	
	Film
Hirst	2 – Preferred to go to the cinema.
West	4 – Would make use of it, especially if it could be directly connected to a TV.
Smith	3 – Would use it, although DVDs currently suffice her needs.
Howarth	4 – Thinks it would be easier to watch films “there’s no need to physically find the film...the extra choice would be welcomed”.
Maule	4 – Liked the idea, far greater choice, especially old films. Also said that it would aid in watching popular/obscure films “which can be difficult to obtain, say from a video shop. With this system they should always be accessible...” Comment: In terms of popular films, broadcast IP technology would certainly be needed to reduce the bandwidth requirements.
	Book (including scientific journals etc.)
Hirst	2 – Complained that reading on the screen causes eye strain: “screens would need to be improved so that they can give a more natural look...”. Although he stated cost would influence his use of this service (“e.g. if the e-book costs £5 and the paper version costs £20, I’d be inclined to access the e-book”).
West	5 – Would use it, especially for magazines and technical manuals.
Smith	3 – Thinks it would be nice should she need it, but she doesn’t feel she needs to obtain information (in the form of a book) via the Internet.
Howarth	4 – Thinks it would be useful, although the eye strain incurred with VDU use would limit his use, “I wouldn’t use it to read a huge document on a computer screen”.
Maule	5 – Would use it, especially for e-journals. The standardised interface would be greatly beneficial – “at the moment there are lots of databases. This is an inconvenience, it would be nice to have a single searchable area – at the moment these databases aren’t typically connected”.
	Music {radio/MP3s etc.}
Hirst	3 – Liked the idea of not being reliant on one client, such as Apple’s iTunes, which forces users to use their client for music. Would rather use the realms concept if the plug-in interface could work with a range of devices/software programs.
West	4 – Would make use of the facility.
Smith	3 – Would use it if it would be as easy as using CDs.
Howarth	4 – Would make use of the facility.
Maule	1 – Not interested in audio. “I prefer to have a disc and cover in my hand”.

	Shopping {direct interface to shops, compare prices, fridge is able to perform many tasks in an automated fashion}
Hirst	4 – Liked the idea of manufacturers being able to group their products independently of the big supermarkets; especially food shopping. Thought that if a manufacturer's shopping realm could be linked to a supermarket shopping realm via standardised interfaces, it could be easier to compare prices etc. "e.g. if I could have a listing of all the products available, and my fridge was able link these with the local supermarkets, it would be easier to compare prices etc."
West	3 – Saw the use, although wasn't personally interested – "my wife takes care of that..." Is generally concerned about "fresh produce" and on line shopping, "I can't pick out the items I want."
Smith	4 – Liked the idea, although expressed concern that "too much planning" may be incurred.
Howarth	5 – Thought there were obvious benefits in a number of diverse ways. Stated that that this type of technology could help with food education, waste control (*just in time food manufacturing*), as well as price comparison and ordering.
Maule	5 – Thought that this would be useful for comparing prices (e.g. dragging Tesco, Adsa shopping realms onto the fridge device). Also thought it could help with searching for particular items and, coupled with RFID, could be useful for identifying products and their best-before-date. Also thought the potential of this particular aspect is vast – "the system could provide for people with allergies, the system could be linked to the items in the fridge/cupboards etc. and provide menus tailored to their needs...perhaps an allergy realm (the functionality) could be dragged onto a fridge to warn of dangers etc."
	Web-pages {information/typically in the public domain, organised according to category}
Hirst	3 – Stated he would be happy to traverse several realms to limit the search domain. Suggested a realm search facility would be useful.
West	5 – stated that this concept would greatly help in minimising the "searchable area". Also stated that realms could be added to a "favourites menu" in a similar way to web-pages.
Smith	2 – largely felt satisfied with current search facilities, such as Google – thought that for simple searches, realms would provide limited use and likened them to the "advance search facility" found on Internet search engines.
Howarth	2 – Felt that current search engines provided reasonably good search results. Expressed concern on the ability of current technology to group web-pages/sites to form realms, citing "enormous amount of web-pages".

Maule	4 – Felt the Realms concept could help limit search results – "anything to reduce the number of hits...this technology could help focus and direct the searcher... Also the concept of 'communities' could be useful, because 'peers' within a realm are likely to be related..." Comments: The aspect of 'communities' has been tackled before, with limited success. The concept of web-rings revolve around a similar philosophy and Google uses an aspect of this to organise search results. As part of the search algorithm, Google compares the number of websites that reference a page/location. Google takes the view that if an article is useful, other authors will acknowledge this and reference/link to it.
	Telephony {information/Address}
Hirst	3 – thought it could provide some benefit and help bridge the gap between current instant messaging programs (user profiles etc.) and traditional phones.
West	5 – Thought that it would be beneficial, especially for businesses. Also stated it would provide an alternative to directory inquiry services.
Smith	2 – Felt content with the current system. Thought that people should simply store contact phone numbers on devices as they are now.
Howarth	4 – Thought there were benefits, especially if someone was to change their contact number etc.
Maule	2 – Not a great telephone user. Not really interested.
	Software {sorted/categorised}
Hirst	4 – Felt this would be very useful, especially for software location and comparison.
West	5 – Thought this would solve many problems associated with software location.
Smith	3 – Didn't feel she would use it because she rarely downloads/installs programs.
Howarth	4 – Thought it would be useful for comparison purposes, and aid the location of software based on the category rather than the abstract name.
Maule	5 – Thought that the concept could enable users to drag software interfaces onto a computer, reducing the difficulties associated with installing and updating software.
	Abstract Family realm {one that a member of the family "creates and configures" for family use, e.g. to hold family records (medical/photos), contact details (e-mail/phone/addresses) etc. Friends could be an aspect of/included in such a realm.}
Hirst	2 – Not sure he would have the time to set up or maintain such a realm.
West	2 – "Some aspects would be useful, although for my needs, a mobile phone is sufficient to store phone numbers etc. I'm not sure what real benefits I'd gain from a virtual existence."

Smith	4 – “I think that would be really good because it would provide a place where all my family could virtually congregate. It could improve the features available on Instant Messaging systems like MSN Messenger, such as leaving messages for people, especially if the family realm could always be available”.
Howarth	3 – “It would be good if a system like this could exist, although I’m not sure if many people would be capable of creating such a thing. I also think that this sort of system would only be useful for people that understand how computers work. Maybe in twenty years time it would be more useful – when the general population have a better understanding.”
Maule	4 – “Yes, I can see the benefits of such a system. It would be lovely to share pictures with other members of the family, especially if it was simple to use. This could be greatly beneficial over current systems...e-mail for example, I have to e-mail specific images to each member...it would be nice to dump them into a realm and just say that the photos from the event are in the following folder and all family members can access them, comment on them etc. I don’t really have time to fiddle with my on line web-space, and it isn’t really possible for the family members to update content easily.”
	Suggested Realms
Hirst	None.
West	Felt this concept could be useful in **genealogy**, providing a public space to collaborate in – similar to Wikipedia in terms of collaboration.
Smith	**Medical** realm with access to on line chemists and doctors.
Howarth	A **public services** realm providing information and an interface to public services, e.g. Waste disposal times and plumber call-outs.
Maule	Felt this concept could be useful in **genealogy** and tracing people. “Nodes could represent humans, and I could associate permissions for accessing telephone numbers, addresses etc.”
5) Are you happy to pay for resources on the Internet, e.g. books, films, music etc.?	
Hirst	Yes. Would not pay for using the realm to search/gain access to the protocol, but happy to pay for a product/service. However, he would demand a lower price because of the lower overheads. Dale thought that this is only fair for electronic resources because there aren’t raw material costs, and the realms concept in most situations removed the “middle men”.
West	Yes, would be happy to pay for any resources he used.
Smith	Yes, although wasn’t sure if she would access the resources often which were discussed in prior questions.
Howarth	Yes, happy to pay for any resource he accesses.
Maule	Yes – not pay for the search, but for the item.
6) If “yes” to question 5, what form of payment would you prefer: Credit/Debit Card, Cheque other?	
Hirst	Credit or debit card.

West	Credit or debit card.
Smith	Credit or debit card.
Howarth	Credit or debit card.
Maule	Credit or debit card.
7) Should a pay-as-you-go/top-up service be introduced, where credit could be bought in a similar way to pay-as-you-go mobile phones. Would you make use of this service over other methods (e.g. Ccs/DCs)?	
Hirst	Would make use of this service instead a debit card. Thought that this could be used as a "credit control" method, e.g. allocate about £10 per week on entertainment.
West	Would not make use of this service, although he stated that he understood its significance.
Smith	Would make use of this service, citing that it acts as a credit control mechanism. Furthermore, she thought that there are fewer security risks with this method.
Howarth	Would make use of this service. He said that he liked to know how much money he has spent. He stated that this reason is part of the appeal of flat-rate fees (e.g. broadband tariffs).
Maule	Would use it to limit spending.
8) Would you be more concerned with having your location details listed in a telecommunications realm? {similar details to that of a public phone book}	
Hirst	Happy with having his details listed, so long as he had control of the content.
West	Was happy with having his details listed and felt it was especially important for business.
Smith	Happy with having her details listed.
Howarth	Is concerned – thought that this type of system could be too pervasive, especially with "address to name" resolution as opposed to "name to address" resolution.
Maule	"In principle I don't really mind...although I am increasingly worried about nuisance phone calls. If this system made it easy for spammers and a flat-rate phone system was widespread, I think I'd be very concerned. Comments: Perhaps an authorisation system to communicate, such as that found in common IM programs, would be required.

Bibliography

Hirst, D., Personal Communication 29 July 2004.

Howarth, R., Personal Communication 30 July 2004.

Maule, A., Personal Communication 30 July 2004.

Preece, J.; Sharp, Rogers, Y.; Sharp, H. (2002) *Interaction Design*, Massachusetts: John Wiley & Sons Inc.

Smith, K., Personal Communication 30 July 2004.

Thompson, M. K. and Weal, M. J., Personal Communication 23 July 2004.

West, T., Personal Communication 29 July 2004.

Section 6: Study of realm interfaces

This section will focus on different aspects of the realm interface and include all types of devices from physical PCs, washing machines etc. to virtual devices such as virtual videos, doctors etc.

6.1 Realm Location

Although this isn't concerned specifically with the realm interface, it is important to study where realms should exist, and thus where the interfaces reside.

It is evident from Section 4 that the realms need to be logically placed. This topic, however, is extremely large and is a research area of its own. Consequently, it has only been possible to provide a summary.

A semantic categorisation system is clearly required; Irwin (2001) provides four principles upon which General Classification Theory is based, see Figure 6.1.

General Classification Theory Principles

1) Every thing, object etc. has to have a distinct and unambiguous description of its unique qualities.

2) Principles involving likeness and distinctness must be used in creating classes.

3) Hierarchies and other relational methods are necessary in order to group fundamental characteristics and to identify fundamental differences clearly.

4) The final system should appear as a logical progression from general to particular.

Figure 6.1 Source: Irwin (2001)

One such system is the long standing Dewey Decimal System, commonly used in libraries to organise books. This breaks down all manner of books into categories, where further divisions are made, making each topic more specific, see Figure 6.2.

This type of system may be sufficient for the organisation of traditional information (*such as an encyclopaedia*). However, in many respects the Realm scope is far greater as it attempts to bridge the gap between traditional information requirements and modern information requirements. That is, modern information requirements are more diverse because of the ability to supply often dynamic information automatically and near instantaneously from different environments.

Consequently, the Dewey Decimal System would need to be amended to account for all manner of different information sources and include aspects such as entertainment and communication. An additional problem with this approach is where to place cross-disciplinary topics since they require a unique location. For example, the Dewey Decimal system places "the mathematics of bell-ringing" under “Fine Arts”.

Further to this and of particular concern, is that according to the Online Computer Library Center (2004), they own the copyright rights to the Dewey Decimal System. This would make this particular system incompatible with Realms because it would be infeasible to pay a license fee per realm. It must be noted that this is only one of a large number of popular classification systems including:

- Library of Congress Classification (LCC)
- Universal Decimal Classification (UDC)
- Colon Classification (CC)
- Bliss Classification (BC)

Each would have to be analysed for their merits and deficiencies. However, one solution to classification is for a committee (*such as the IETF*) to create top level realms and allow the "Internet community" to freely form and alter sub-realms. This could provide a "survival of the fittest" model, where realms are located according to where the majority of users access them. This could be achieved by initially creating a number of symbolic links to the ontology, where links (*nodes*) that are rarely used are removed over time.

It is evident that this is a complex matter and that further research is needed to determine the most appropriate realm classification system.

Dewey Decimal System: Categories to traverse in order to obtain the Butterflies category

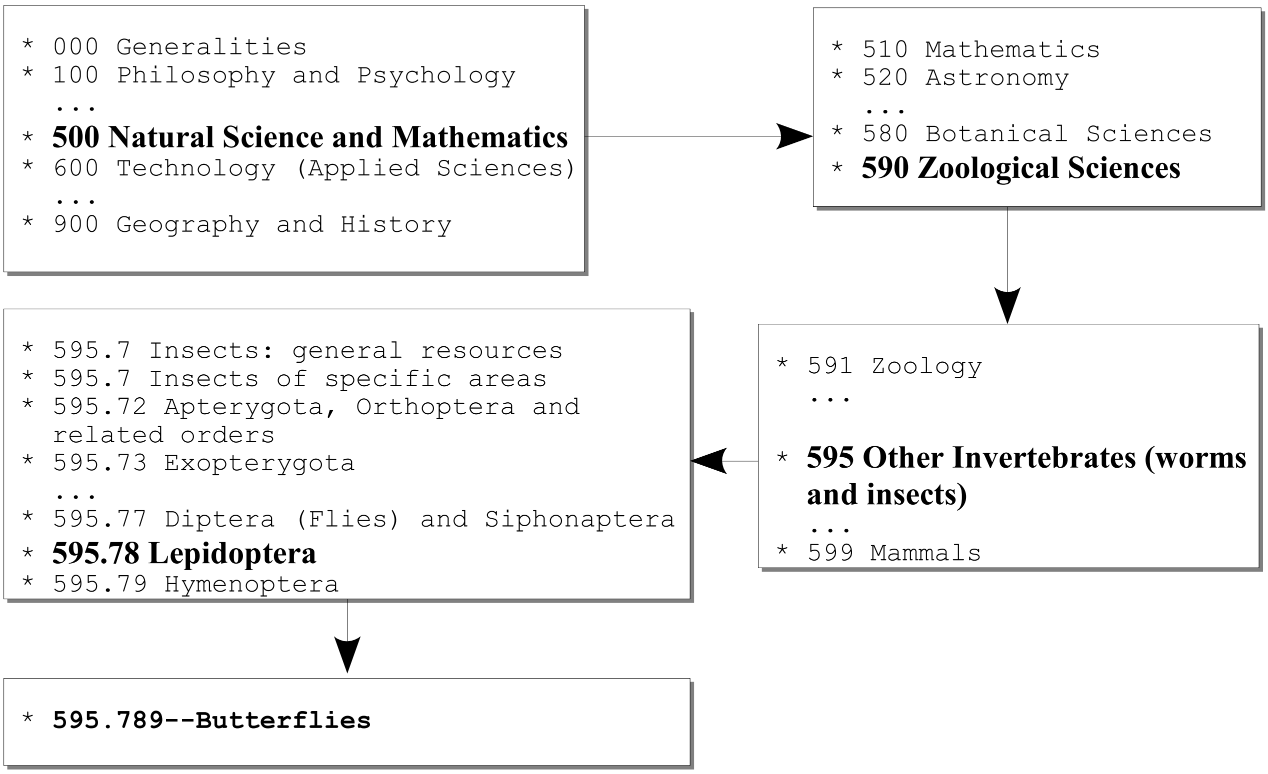

Figure 6.2

Adapted from: Vesper (2004)

6.2 Connection of devices

The Realms concept is founded upon peer-to-peer (P2P) computing, where all nodes (*known as peers in P2P*) have equal status within the network (*in this case, a realm*). As more devices connect onto the realm, the network's capacity increases in bandwidth, processing and storage etc. and there is no additional strain on a central point because P2P is a distributed network. Thus there isn't a central point. However, despite the benefits of scalability over the traditional client server architecture, P2P is prone to complexities in network management because there isn't a central point to administer backups and set security policies etc. Consequently, P2P has only relatively recently been a practical network architecture.

This infancy has resulted in a number of different protocols and implementations. JXTA (pronounced *jux-ta*) however, aims to solve the lack of standardisation by providing an Open Source P2P computing platform where P2P applications can be built, see Figure 6.3 for the JXTA abstraction and Figure 6.4 for the network structure. The project aims to support all manner of devices ranging from mainframes to sensors and is focused on an architecture that is independent of programming language, platform and network, see JXTA (2004a) for full details.

It is evident from JXTA (2004b) that the system comprises of six protocols (*shown in Fig. 6.5*). According to the authors, these are specifically designed for ad hoc, pervasive, and multi-hop P2P network computing. Further to this, the protocols enable devices to form self-organised and self-configured peer groups which are network location independent (*edges, firewalls etc.*) and don't require a centralised management infrastructure. This means that JXTA could provide a solid P2P foundation for a Realm infrastructure. Moreover, because JXTA is an Open Source "community effort", a framework has been created (*and continues to be developed*), that enables standard developers to produce results (*systems*) that they typically could not achieve without the community support.

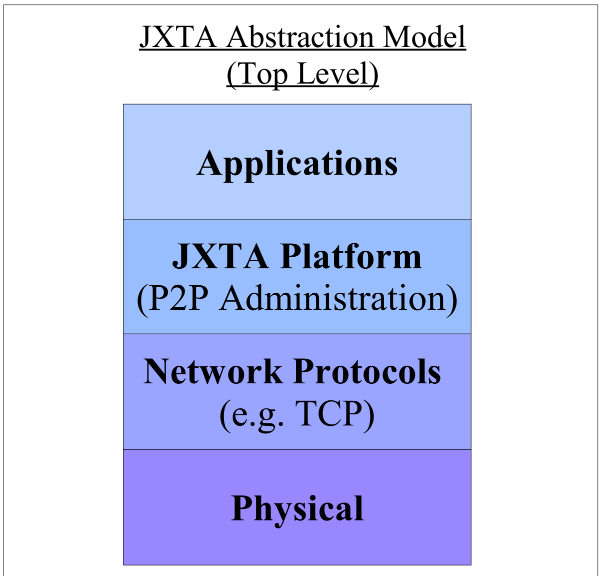

Figure 6.3

JXTA Virtual Network Structure

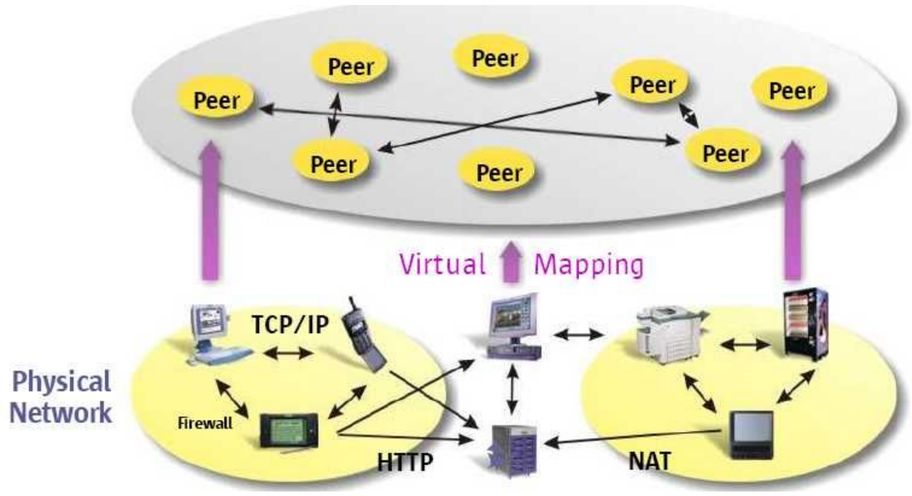

Figure 6.4

Source: Sun Microsystems (2004a, p.5)

JXTA Protocol Stack

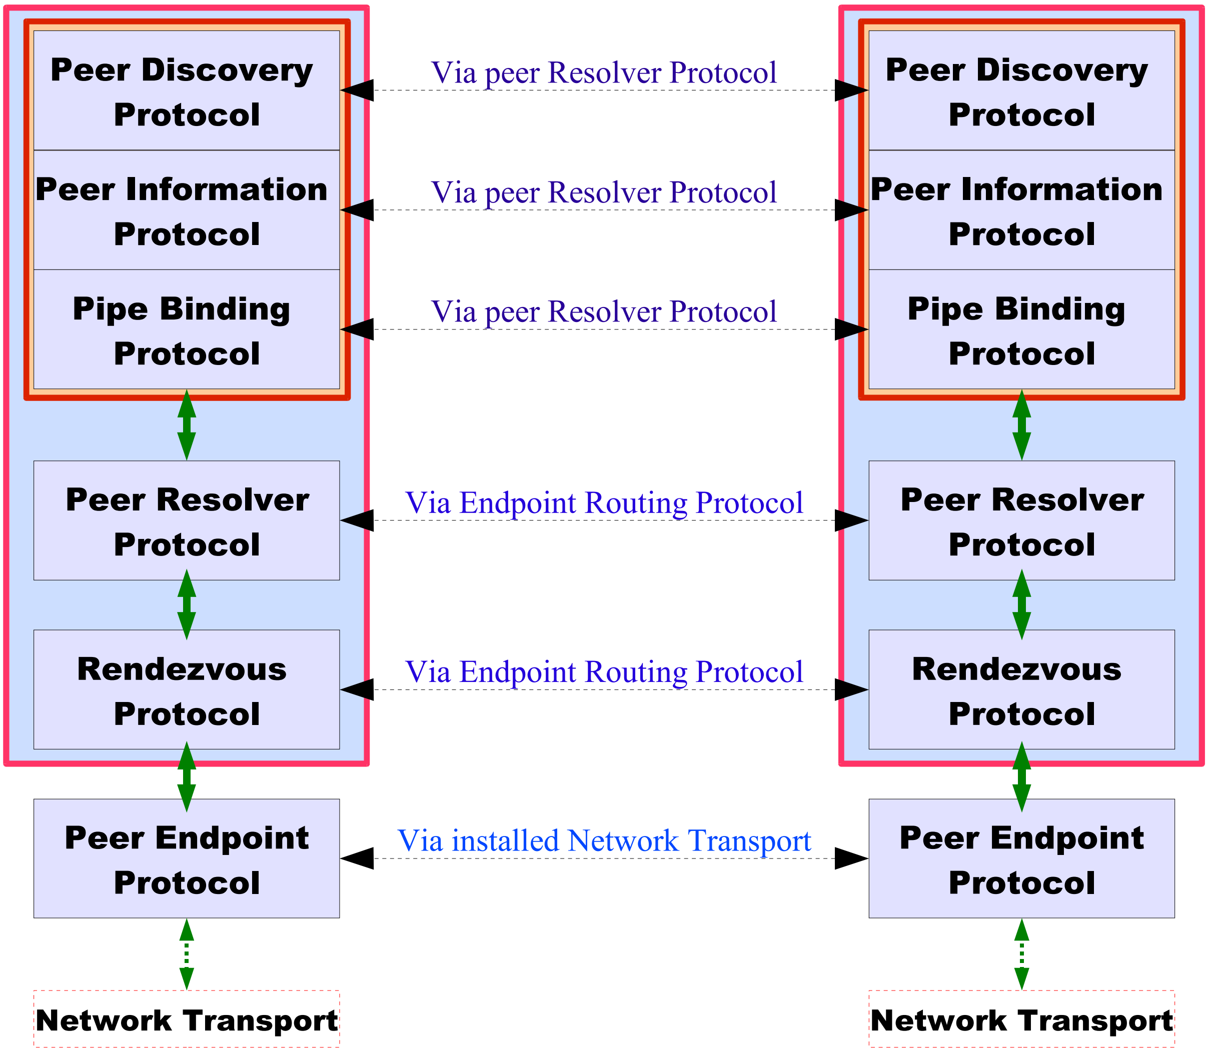

Note: the endpoint protocol can be plugged into the stack for different networks:

* **TCP/IP** to Internet or Intranet
* **HTTP** through firewalls
* **SMS** to cellphone network
* **Packet radio** to packet network
* **Other customer 'driver'** to other network

Li (2002) states that an endpoint protocol's only requirement is to be able to send or receive XML messages between two physical nodes. There is no requirement for reliability in transmission or message broadcast.

This means that simple and complex protocols can exist at this level. This explains why HTTP and TCP/IP are found at the same layer within the JXTA protocol stack.

Figure 6.5

Image adapted from: Taylor, I. (2004, p.18)

6.3 JXTA and Realm Services

As discussed throughout this document, the entire Realms concept is based upon the fact that nodes, with a subject in common, unite to create a realm which provides services that enable node location (*within the scope of the realm*) and enforce realm standards regarding interface (*i.e. node attributes, methods and interface codecs*) that the realm owners (*either private or community representatives*) are able to specify.

From this perspective, the system appears to act in the same way as JXTA ontologies that have been ordered to suit the conceptual needs of the human users. However, it must be noted that JXTA provides a tool to abstract the underlying network architecture and support applications such as a realm. Thus JXTA and Realms are at two different and distinct levels of abstraction, where realms are primarily at the Application level, see Figure 6.6.

JXTA Search Network

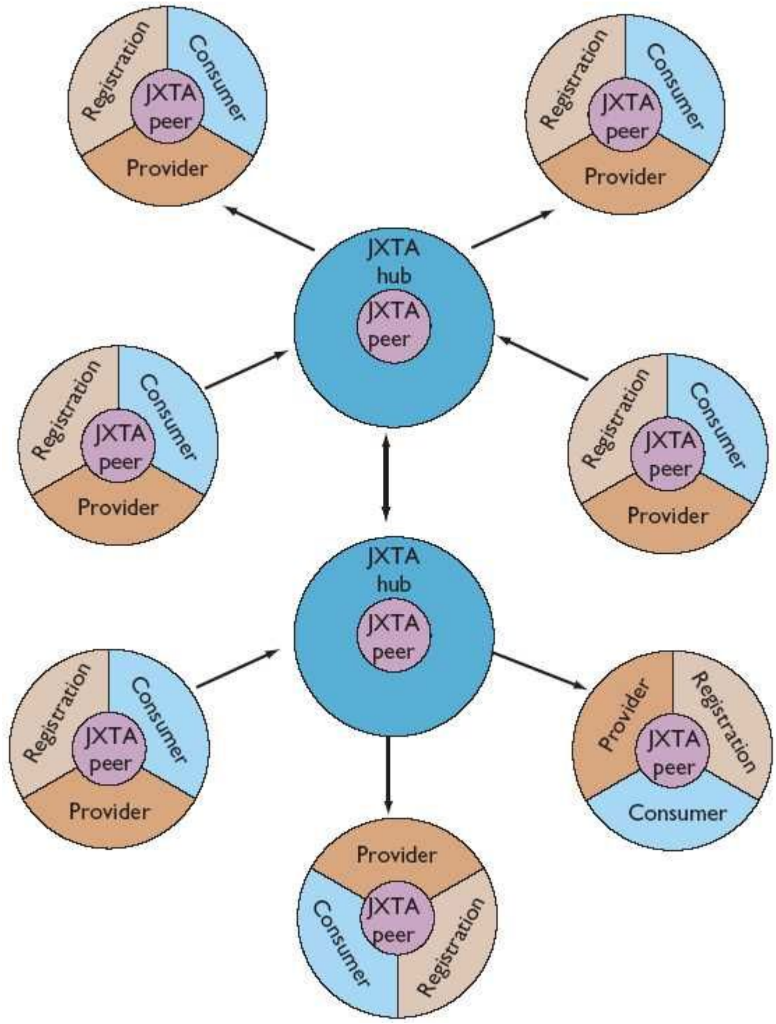

Figure 6.6 Source: Waterhouse *et al.* (2002, p.2)

Further to this, the Realms concept is focused on top-level human and computer system aspects:

- Context – realms and nodes have been given a human understandable context.
- Leadership – realms and nodes have been given human leadership, something that in terms of human group management is vital for managing groups, see Guirdham (2002) for details on human leadership. Parts of this may be automated, such as voting systems. However at this high level, ontologies are under human control regardless of the true underlying formations.
- Functionality – realms provide high-level functionality to humans and computer systems.

It is evident from the literature that although JXTA could support the underlying infrastructure, the concept of supplying the same functionality by different means within an ontology is not inherently supported. For example:

E.g.

Companies A, B and C offer a JPG generation service. 'Company A' uses the ABC algorithm, 'Company B' uses the FOO algorithm, and 'Company C' uses the XYZ algorithm. All companies adhere to the same protocols (*attributes, methods, codecs etc.*).

This is not possible because it's not the purpose of the JXTA system. As described earlier, JXTA is a platform that abstracts the underlying architecture and is not a realm application.

In essence, the basic JXTA system could:

- administer all node data in a P2P fashion (*robust node storage*),
- assist in node location (*robust node location*),
- enable a wide variety of devices to seamlessly communicate,

and provide appropriate mechanisms to support:

- centralised control,
- a collaboration layer which specifies the nodes interface (*attributes, methods, codecs etc.*). Sun Microsystems (2004b) state that Peer and PeerGroup services "*describe not only the service being provided, but also where to obtain the specifications and executable code for the service*".
- individual implementations of a service, *e.g. the various JPG generation services in the previous example*,
- security, it could be possible for the ontology leader to digitally sign elements.

Figure 6.6 provides an overview of the JXTA search framework. It is evident that peers can be both providers and consumers of information and that they can form hubs which are able to efficiently route queries from consumers to information providers.

Klampanos (2002) points out that it is possible for hubs to be chained together, implying that a query can be routed to more than one hub and therefore to more than one information provider. This aspect is also shown in Fig. 6.6.

Summary

JXTA supports very dynamic networks and provides a significant grounding which many P2P networks could benefit from. If the Realms concept were to be seriously considered for implementation, then the JXTA system should certainly be considered as both concepts share many principle ideologies. Further to this, JXTA caters for, but is by no means limited to, XML specified protocols. This helps to keep protocols within the system lightweight and easily portable. JXTA can support a range of devices ranging from mainframes to sensors.

An important point to raise is that Realms, as described throughout this documentation, are a higher level abstraction than JXTA ontologies. To help illustrate this, one might try comparing a JXTA hub to a realm system, merely because it specialises in a specific area. This is not the case as a hub is primarily for computational optimisation (*used to locate data*) as opposed to a human abstraction tool.

Bibliography

Denny, M. (2002) *Ontology Building: A Survey of Editing Tools*, California, O'Reilly Media Inc.
Available from: http://www.xml.com/pub/a/2002/11/06/ontologies.html, [Date Accessed: 2 August 2004]

ERCIM NEWS (2002) *European Research Consortium for Informatics and Mathematics Special: Semantic Web*, France, ERCIM.
Available from: http://www.ercim.org/publication/Ercim_News/enw51/EN51.pdf, [Date Accessed: 2 August 2004]

Gärdenfors, P. (1995) *Human communications - what happens?*, Lund University, Sweden.
Available from: http://www.lucs.lu.se/People/Peter.Gardenfors/Articles/Communication.html, [Date Accessed: 8 August 2004]

Guirdham, M. (2002) *Interactive Behaviour at Work* (3rd edition) London: Prentice Hall.

Irwin, S. (2001) *Classification Theory and the Internet: A move toward Mulitdimensional Classification*, Colorado, University of Denver.
Available from: http://www.du.edu/lisa/students/student_papers/Irwin.htm, [Date Accessed: 3 August 2004]

JXTA (2004) JXTA v2.0 Protocols Specification, California, Sun Microsystems.
Available from: http://spec.jxta.org/nonav/v1.0/docbook/JXTAProtocols.html#id2796266, [Date Accessed: 4 August 2004]

Kelaskar, M. (2004) *Project JXTA - Peer Empowerment*, New Jersey, The State University of New Jersey.
Available from: http://paul.rutgers.edu/~kelaskar/JXTA.ppt, [Date Accessed: 3 August 2004]

Klampanos, I. A. (2002) *JXTA Search*, Scotland, University of Glasgow.
Available from: http://www.dcs.gla.ac.uk/~iraklis/fyp_report/node25.html, [Date Accessed: 8 August 2004]

Krishnan, N. (2001) *The Jxta solution to P2P*, California, JavaWorld.com.
Available from: http://www.javaworld.com/javaworld/jw-10-2001/jw-1019-jxta.html, [Date Accessed: 4 August 2004]

Li, S. (2002) *Making P2P interoperable: Creating JXTA systems*, New York, IBM.
Available from:
http://www-106.ibm.com/developerworks/java/library/j-p2pint3/, [Date Accessed: 6 August 2004]

Malcolm, C. (2003) *Machines with Minds?*, Edinburgh, University of Edinburgh.
Available from: http://www.dai.ed.ac.uk/homes/cam/IAS_docs/Robot_Phil.shtml, [Date Accessed: 8 August 2004]

Online Computer Library Center (2004) *Dewey services [OCLC - Cataloging and Metadata]*, Ohio, OCLC Online Computer Library Center.
Available from: http://www.oclc.org/dewey/default.htm, [Date Accessed: 2 August 2004]

Protégé Project (2004a) *The Protégé Ontology Editor and Knowledge Acquisition System*, California, Stanford Medical Informatics.
Available from: http://protege.stanford.edu/, [Date Accessed: 3 August 2004]

Protégé Project (2004b) *Protégé OWL Plugin - Ontology Editor for the Semantic Web*, California, Stanford Medical Informatics.
Available from: http://protege.stanford.edu/plugins/owl/, [Date Accessed: 3 August 2004]

Skype Technologies (2004) *Skype*, London, Skype Technologies S.A.
Available from: http://www.skype.com/, [Date Accessed: 4 August 2004]

Sterner, W. H. (1997) *The Loss of Meaning in Human-Computer Communication*, Chicago, University of Chicago. Available from: http://www.billsterner.cs.uchicago.edu/eliza.pdf, [Date Accessed: 8 August 2004]

Sun Microsystems (2004a) *JXTA Technology: Creating Connected Communities*, California, Sun Microsystems Inc. Available from: http://www.jxta.org/project/www/docs/JXTA-Exec-Brief.pdf, [Date Accessed: 6 August 2004]

Sun Microsystems (2004b) *JXTA: P2P Grows Up*, California, Sun Microsystems Inc. Available from: http://java.sun.com/developer/technicalArticles/Networking/jxta/, [Date Accessed: 8 August 2004]

Taylor, I. (2004) *JXTA Lecture*, Cardiff, Cardiff University. Available from: http://www.cs.cf.ac.uk/user/I.J.Taylor/DistributedSystems/PPTSlides/JXTALecture.ppt, [Date Accessed: 8 August 2004]

Vesper, V. (2004) *Dewey services [OCLC - Cataloging and Metadata]*, Ohio, Middle Tennessee State University. Available from: http://www.mtsu.edu/~vvesper/dewey.html, [Date Accessed: 2 August 2004]

Waterhouse, S., Doolin, D.M., Kan, G., Faybishenko, A., (2002) Distributed search in P2P networks, *IEEE Internet Computing*, 6(1), 68-72.

Section 7: Critical Analysis of the Realm concept for the Internet structure

7.1 Realm recapitulation with respect to the Internet structure

A significant aspect of the Realms concept is that it attempts to give nodes a human understandable context. The purpose is to help increase the ability of a human to locate a node, be it a physical or virtual device (*including web pages and information*). The perceived requirement stems from the fact that many of the current systems are keyword based (*which often neglect the human context*) or are dependent on a user "browsing" in an unguided manner. This is typically the case with current search engines such as Google, Yahoo, Altavista, Lycos and the like. These search engines do offer limited categorisation services. This is discussed later in this section.

Another important aspect is that the Realm concept encourages standardisation between nodes in so-called "common ground". An analogy: *Speakers from around the world attend a global conference. They all speak their mother tongue but they choose to converse in English because this is the accepted protocol for international communication*. This same concept could be applied to all manner of ontologies for primary protocols (*such as text, audio and video*) and high-level information protocols such as RSS* which is used to syndicate news content with the option to provide broadcast updates to subscribers of the feed.

7.2 Critique of Internet content categorisation

First, many other Internet services provide content categorisation. However, the categorisation facilities are limited to web pages, therefore devices (*such as IP Phone directories*) aren't indexed. Further to this, the search engines are in the private sector, and therefore the proprietors have overall control. They reserve the right to charge for listings, list only the content that they desire, and control any development of a categorisation system (*i.e. the development on the system is controlled by a private organisation and not by anyone wishing to improve the service*).

Figure 7.1 provides a screen shot of Google's Index and Directory web pages. One point to note is that the directory section of Google isn't immediately obvious from the index page. Consequently, we thought it would be interesting to conduct a rudimentary poll from regular Google users, to establish what fraction actually knew, let alone used the service. The results are shown in an appendix to this section.

It is evident, from the results shown in that appendix, that 80% of the users weren't aware of Google's directory facility, although 60% had used categorisation facilities in the past. Further research, with a greater number of users, is needed to determine the exact causes, although one can speculate that education and usability are significant factors.

Chen *et al.* (1995) state, with regard to categorisation, that "*browsing allows users to explore only a very small portion of the large Internet information space*". This is a possible explanation for why some users didn't find content categorisation particularly useful; some found it time consuming or difficult to locate the required content.

* The authors of RSS have never provided a firm definition, so the following definitions are all used interchangeably to describe the same concept: Really Simple Syndication, RDF Site Summary, Rich Site Summary

Google Directory Service

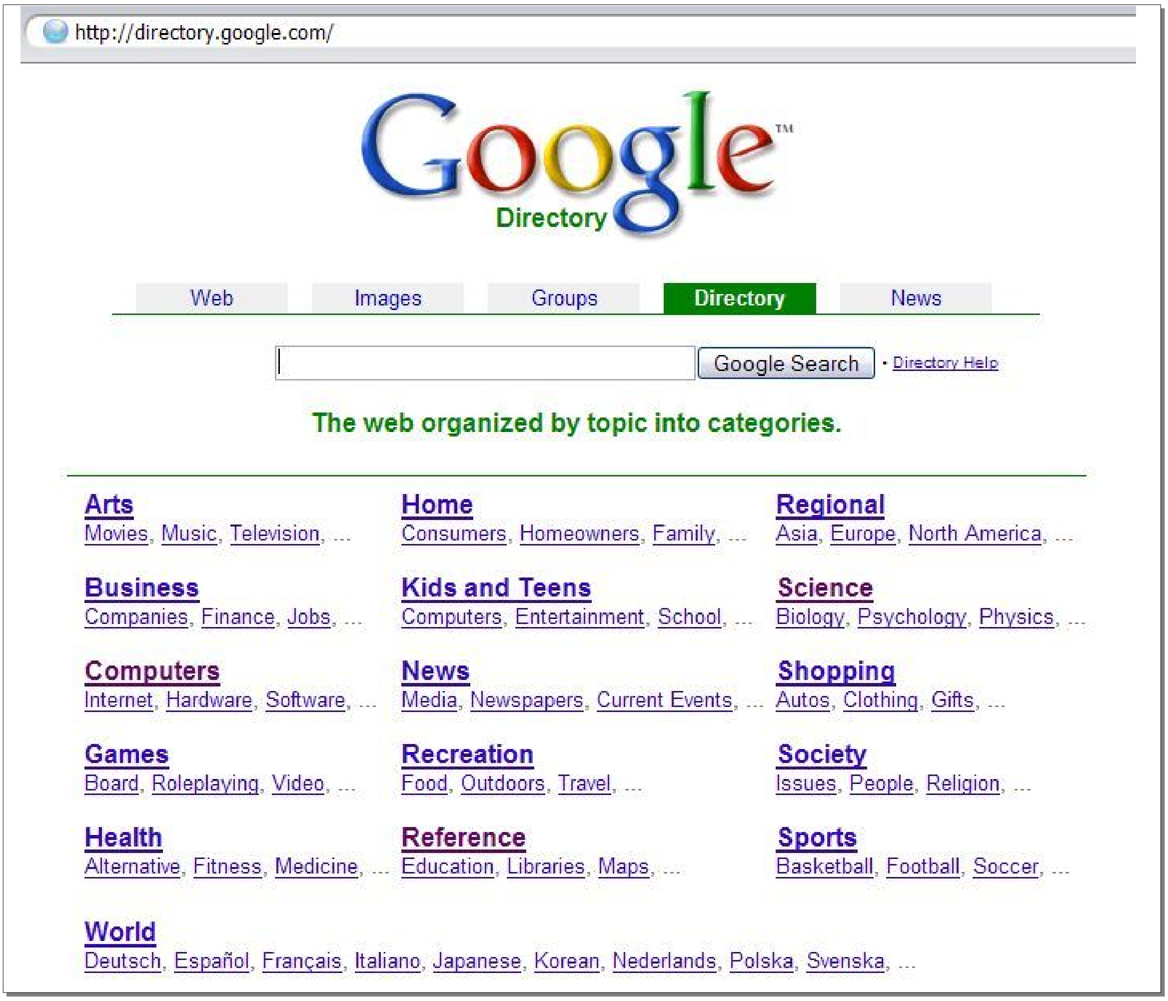

A — Source: Google (2004a)

Google Directory Service

B — Source: Google (2004b)

Figure 7.1

An e-mail was sent to Google (see Appendix) which was focused on content categorisation. It was hoped they could provide some valuable insight. However no reply has yet been received from Google.

In addition to the fact that many search engine websites provide category organisation, Chen *et al.* (1995) highlight problems with this approach.

7.2.1 Vocabulary

Chen *et al.* (1995) state that this is a consequence of diversity of expertise and backgrounds of system users. Further to this, they state that there is a prevailing inability of searchers themselves to fully articulate their needs. In relation to Realms, the questions of "who decides what a realm should be called" and "where it should be located" are important. To propose an exact solution to this for the Realms concept would require further research. However, one can hypothesise that symbolic links (*such as those available on the Unix operating system*) could play an active role. The following is an illustration of this.

E.g.

Symbolic links in the Unix file system act as pointers to the true location. Thus a directory or file can be linked in the following ways:
- the symbolic directory "/easyDirectory" could be linked to "/hard/to/remember/directory/"
- the symbolic file "/easyFile.txt" could be linked to "/hard/to/remember/file/location/File.txt"

In both these cases the symbolic links enable seamless redirection. This concept could be applied to Realm names, where synonyms and foreign equivalents (*e.g. English: Dog and German: Hund*) point to the same nodes/ontologies. However, it must be noted that some authors, such as Metcalfe (2000, p.59), have cited a long-term worldwide trend towards fewer languages, which may have an affect on this complexity.

Further to the problem of synonyms, Chen *et al.* raise the problem with the semantics of a word. It is evident from their paper that a concept may be perceived differently by different searchers and that it may also convey different meanings at different times. The examples below help explain the extent of the problem.

- Deprecated Terminology (*also deprecated semantics*)
  Frenkel (1994) states that in fields such as biology, scientists' knowledge of a concept under study often changes (*relative to time*). He states that this causes problems with vocabulary as scientists unaware of the developments use deprecated terminology.

- Incorrect Terminology
  Incorrect terminology is often used, especially by the inexperienced. For example, it is common, in some circles, to refer to a USB WiFi adaptor as an "Wireless dongle". However, this term is ambiguous and incorrect in other circles as the word "dongle" refers to a security device that is attached to a computer port for the purpose of preventing unauthorised access to a software program.

- Different Contexts

The word “root” may be interpreted differently according to context. For example, it may be interpreted in botany, computing, linguistics, and no doubt many other contexts.

Further to this, the semantics of a word can be interpreted differently within a context. For example, the semantics of the word root are different within the computing field. It can refer to:

- the superuser of computer system
- the highest level in the directory
- and many other areas.

Moreover, the fact that a user uses the keywords “root” and “computing” does not simply imply that the word “root” is used in the computing sense.

In addition to the vocabulary problems, Chen identifies the following problem.

7.2.2 Information Overload

Chen (1994) states “*the information overload problem is complicated by the diversity of subject area knowledge, classification knowledge, and system knowledge exhibited in different users*”. This, he explains, means that simple browsing systems can potentially confuse and disorient a user, and can cause them to spend a long time browsing whilst not learning anything in particular, coined as the “art museum phenomenon”.

In spite of these problems, there are benefits to be gained from grouping, especially in terms of discovery. For example, once a relevant category is identified, it could be easier to locate relevant information. This helps explain the positive response from a participant, who stated it was beneficial for him to search within a certain category.

It is our view that categorisation and traditional search mechanisms complement one another. Google, for example, provides search facilities within a category. However, the categorisation functionality could be combined with their keyword search to automatically recommend categories, thus reducing the time required for the user to locate the category and empowering them to make a decision. See Chen *et al.* (1995) for more detailed work on Internet Categorisation and Search.

7.3 Critique of Realm Standardisation

The topic of standardisation, especially in terms of the affect on innovation, is a vast subject and has been the focus of many researchers. There are many benefits to be gained from standardisation, such as interoperability and other aspects discussed in this document. However, there are a number of potential adverse affects, especially in terms of innovation. It has been assumed that realms could accommodate innovation by facilitating the creation of sub-realms, private ontologies to extend an ontology’s functionality. However, the “community” approach may cause undesirable effects, such as fallout and potentially unfair practice amongst peers.

In January 2004, the Opensource XFree86 project was “forked” after a heated debate over licensing issues. There are now two projects, X.Org and XFree86.

The standardisation topic is too complex for the scope of this document and consequently won’t be discussed further. However, Feng (2003), Krechmer (2003) and Schoechle (2003) provide detailed analysis of the impact of standardisation on innovation within the IT and telecommunications sectors.

7.4 Critique of P2P to support Realms

Cox *et al*. (2004) provide an insightful paper on DHash, a peer-to-peer distributed hash table based on Chord, see Kaashoek *et al.* (2004), which provides the verification of DNS records. This type of hierarchical structure is envisaged to be similar to the one employed for the Realm structure.

The significant benefit to be gained from the DHash approach is that the concept of a root server is eliminated since the structure is P2P based. This is important since studies shave shown that up to 18% of DNS traffic is destined for the root servers, see Jung *et al.* (2002). This may be partly explained by the fact that Name Servers (NSs) are notoriously difficult to configure. In the study by Jung, 35% of DNS queries never received an answer or received a negative answer, which was attributed largely to improperly configured NSs.

The DHash results, shown by Cox *et al.* (2004), look promising as the system was able to cope well with the dynamic reallocation of data blocks, especially for popular domain records. In simple terms, when a data chunk is in high demand, the system detects this and makes duplicates to cope with demand. However, they noted that the DHash system had additional latency over conventional DNS. Whilst this is a significant setback, especially when minimised latency is important, P2P systems exhibit other benefits. An implementation of Chord is available in JXTA, see Sun Microsystems (2004).

P2P advantages:

- No central point of failure – P2P systems can take advantage of redundancy and load balancing techniques. Chord typically has 6 replicas , see Cox (2004, p.2), which means a significant number of hosts need to be lost before the network experiences loss of data. This makes the systems robust against denial-of-service attacks.
- Efficiency – Efficiency can be improved from node auto-configuration. In the study by Jung *et al.*, 35% of traffic was attributed largely to improperly configured NSs. Cox *et al.* believe that this would be greatly reduced by a P2P system such as DHash.
- Scalability – P2P systems can take advantage of vast amounts of resources: data, computer hardware, and bandwidth.
    - Content based addressing – data can be segregated into specialist grouping based on content. This can create a more refined data repository.
    - Flexible – P2P networks are easy to extend as desired, largely because there isn't a single point of authority for users to ask permission.
- No single owner – a single organisation doesn't need to take responsibility for the management of the whole network. Organisations are able to join and leave as desired.
- Reduced costs – costs inherent with central management are all but diminished. Further to this, the creation of new P2P networks are negligible.

Additional P2P disadvantages:

- Decentralised coordination:
    - It is difficult to maintain a consistent state. P2P systems are typically highly dynamic which can influence the ability to provide services.

    - Protocol management can be difficult with P2P systems because each node is required to operate any given protocol. Different protocol versions and updates pose difficult requirements to the P2P system designers.
    - A distributed service, such as DHash, requires people publishing names to rely on other people's servers to serve those names. This can be a problem in P2P networks, since demand may outstrip the available supply, especially if there is little incentive to run a P2P server. This is often evident in P2P MP3 file sharing systems, where demand often exceeds the available supply, resulting in slow downloads.
- Lack of centralised control:
    - are inherently difficult for authorities to moderate and control,
    - can be taken advantage of by hackers, terrorists etc.,
    - empower users to infringe copyright laws.

- All nodes are not created equal.
    - Computational power and bandwidth have an impact on a node's performance within a network.

7.5 Is there really an alternative approach to pure P2P?

According to UDDI.org (2002), Universal Description, Discovery and Integration (UDDI) "*is a specification for distributed Web-based information registries of Web services.*"* The purpose is to provide programmable elements which others (*primarily B2B*) can access. For example, assume a courier provided a cost calculation Web Service. This functionality could be inherited by other systems, for example e-commerce businesses. The purpose of the UDDI framework is to publish Web Service information pertaining to any given company. This provides an automated system that relieves the need for rudimentary communication, such as phone calls and e-mails, to provide business and Web Service information. Further to this, it provides a mechanism to track changes and ensure consistency in the service description syntax. Business information, including their service descriptions, are implemented in XML.

UDDI is perceived by its authors as the next layer in an emerging stack enabling rich Web services

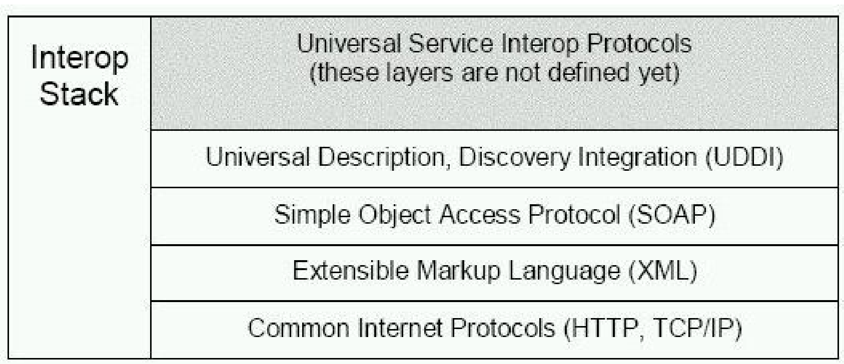

Figure 7.2 Source: UDDI.org (2002, p.4)

*A web service "*describes specific business functionality exposed by a company, usually through an Internet connection, for the purpose of providing a way for another company or software program to use the service*". UDDI.org (2002, p.3)

UDDI.org (2002, p.4) states that the UDDI specifications borrow the lessons learned from XML and SOAP to define the next layer, enabling two companies to share a means to describe their services and query each other's capabilities. Figure 7.2 illustrates the UDDI in relation to its underlying protocol stack.

At a high level, the UDDI framework is fairly simplistic. Central registries store UDDI data in a similar way to the DNS; distributed client server nodes. This structure is hierarchical as opposed to fully distributed, as with P2P. There are four aspects to the registered data.

Businesses register public information about themselves:

- **White Pages** – contain the business name, text description, address, contact information, and other related information.
- **Yellow Pages** – contain the businesses classification and the services it offers, for example: a company may manufacture kayaks and provide cleaning services.
- **Green Pages** – contain information about how to invoke the offered services.

Standards bodies, programmers, and businesses information about Service Types:

- **Service Type Registrations** – contain details pertaining to a certain type of service.

Summarised from: UDDI.org (2002)

The UDDI framework described, provides a conceptual structure that is comparable to that of realms. Figure 7.3 is an illustration of the conceptual UDDI structure.

Illustration showing the relationship between the technical discovery layers defined by UDDI and the search services

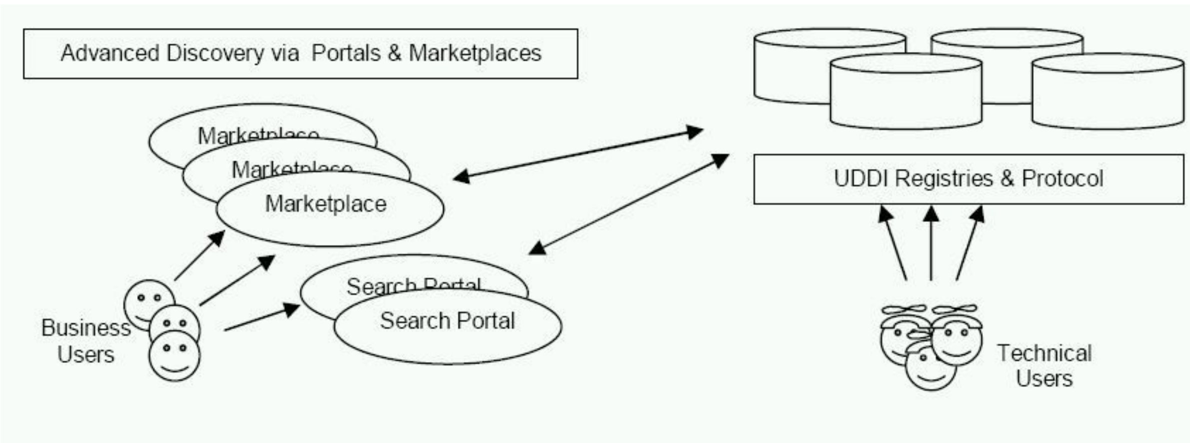

Figure 7.3

Source: UDDI.org (2002, p.6)

The UDDI approach solves a number of problems that are intrinsic with P2P systems. Primarily, centralised administration enables the authorities and Internet consortiums to maintain some control. However, there are potential disadvantages when compared to P2P systems; Denial of Service attacks leave systems vulnerable, and network connections may not be used efficiently as the network is focused on single sources, e.g. a server.

It is evident that both P2P and Client Server computing each have their benefits and disadvantages. Client Server has been used for many years and experiences the benefits from being a tried and tested architecture. Conversely, P2P is a relatively new practicable architecture and inherits the associated problems.

Laoveerakul *et al.* (2002) provide a paper on decentralised UDDI and have provided a system which enables UDDI to work on a P2P architecture.

It is possible to have a hybrid architecture. For example the first MP3 Napster client used centralised servers to list content whilst it enabled peers to transfer directly between each other.

Appendix

Categorisation Questionnaire

Please indicate your computer education level. (e.g. IT Professional, including CS degree or equivalent, Power User (PC builder/OS Installer), Standard User.	***Are you a regular search engine user? (e.g. Google, Yahoo etc.) - If so, please indicate your usual search engine.***
Pau (2004): IT Professional	Pau: Yes – Google
Snowden (2004): Power User	Snowden: Yes – Google
YoOng (2004): IT Professional	YoOng: Yes – Google
Liu (2004): IT Professional	Liu: Yes – Google
Smith (2004): Standard User	Smith: Yes – Google
Maule (2004): Standard User	Maule: Yes – Google

Have you heard of the Google Directory service, used to categorise the content available on Google? (directory.google.com)	***If you have used another website categoriser, what website did you use (e.g. Altavista, MSN, Lycos, other etc.). Please comment.***
Pau: Yes	Pau: Yes – used Yahoo and Google's directory. I don't find them that useful.
Snowden: No	Snowden: No.
YoOng: No	YoOng: Yes – used Yahoo, limiting the search criteria to only computers. Can be time consuming, but can be worth the effort, especially when looking for related material.
Liu: No	Liu: Yes – Yahoo. When I first used Yahoo, I went through the categories – it seemed most logical. I then changed to keyword search because the categorisation results didn't always contain the required content.
Smith: No	Smith: Yahoo – used it a few times, although I didn't find it that useful. It was difficult to locate the required content through the categorisation system.
Maule: No	Maule: No – wasn't obvious.

If you weren't aware of the categorisation services, would you use them now? If you have or do use them, please comment on your usage patterns with these services.
Pau: "I have used them, although I think that they take too long to traverse. Also, when I search, I don't necessarily want to exclude other categories as this might mean I'm not aware of certain information."
Snowden: "Yes I would use them, although I'm not sure I'd use them for every search because it doesn't look like it would be as quick as the straight search for basic/very specific searches."
YoOng: "I was aware of these services. Their use depends on what I'm searching. If I find it's difficult to locate the required information, then I use this service."
Liu: "I was aware of the categorisation service, but I haven't tended to use it because it's typically too slow, although I would use it if it was possible to search and click recommendations which restrict the results". e.g. search for ".NET" and Google recommends restricting the results to "Computer only content".
Smith: "I was aware of the Yahoo categorisation service, but as stated previously, I found that the system wasn't particularly user-friendly. The keyword search system, although not perfect, does tend to satisfy most of my needs. The advantages of this is that I can locate what I want quickly (e.g. by just typing in one or two keywords) whereas Yahoo's categorisation system requires several clicks, requires me to look at all the categories and typically doesn't result in the most appropriate website."

Maule: Yes – will now use it more regularly and instead of the main search page.

E-mail to Google Inc.

To: directory-feedback@google.com
Subject: Google directory questions

Dear Sir,
I'm an undergraduate research student tasked with the exploration of the Internet as a P2P ontology based system. As part of this research, I am studying information categorisation and I am particularly interested in your directory service.
I'd be grateful if you could answer the following questions:

> Why don't you provide a link to the directory service on your main search page? It's available from "more" --> "Directory". {is this due to popularity and simplicity reasons?}

> Have you considered creating an education zone to show users, not only how to perform basic searches, but how the more advanced systems should be used, e.g. a video/flash animation to show users how to use the directory and other services? As part of my research, the overwhelming majority users (*including experienced users*) weren't aware of Google's directory service, despite being regular users of your service.

> Have you considered allowing the Internet community to create their own ontologies (i.e. groupings in your directory)? {and possibly moderated by community members}

> Have you considered a public voting system to remove inappropriate links? Although your system is pretty good, inappropriate content is often listed – especially sales websites containing inappropriate keywords.

> It is evident that your service specialises in Internet information search (mostly web pages etc.). However, is it foreseen that Google will expand this into other aspects? For example, IP Phones, Electronic Fridges, Sensors etc., are likely to become popular in years to come – does Google plan to expand its search facility into these areas?

> And finally {only applicable for the above question}, has Google considered playing a part in standardisation, e.g. creating ontologies (such as those listed in your directory service) that facilitate common "methods" (similar to programming). An example would be ontology.node.getAuthor(), or ontology.node.playFilm().

If you have answered the above questions, thank you very much, your time is greatly appreciated.

Oliver FJ Snowden.

School of Electronics and Computer Science
University of Southampton
SO17 1BJ
United Kingdom

Bibliography

Chandrasekaran, S. (2002) *Web Services*, Georgia, The University of Georgia.
Available from: http://chief.cs.uga.edu/~jam/home/theses/senthil_thesis/ClassPresFinal.ppt, [Date Accessed: 27 August 2004]

Chen, H. (1994) *An Algorithmic Approach to Building Concept Space for a Scientific Community*, Arizona, The University of Arizona.
Available from: http://ai.bpa.arizona.edu/papers/hicss27w/hicss27w.html, [Date Accessed: 12 August 2004]

Chen, H., Schuffels, C., Orwig, R. (1995) *Internet Categorization and Search: A Self-Organizing Approach*, Arizona, University of Arizona.
Available from: http://ai.bpa.arizona.edu/papers/som95/som95.html, [Date Accessed: 6 August 2004]

Cox, R., Muthitacharoen, A., Morris, R. T. (2004) *Serving DNS using a Peer-to-Peer Lookup Service*, Massachusetts, Parallel & Distributed Operating Systems Group.
Available from: http://www.pdos.lcs.mit.edu/chord/papers/ddns.pdf, [Date Accessed: 14 August 2004]

Feng, P. (2003) *Studying standardization: a review of the literature: Proceedings of 3rd IEEE Conference on Standardization and Innovation in Information Technology held at Delft University of Technology.* Nederlands, Delft University of Technology. p. 99-112.

Frenkel, K. A. (1991) The human genome project and informatics, *Communications of the ACM*, 34(11):41-51. IN Chen (1994)

Google (2004a) *Google Directory*, California, Google Inc.
Available from: http://directory.google.com/, [Date Accessed: 6 August 2004]

Google (2004b) *Google*, California, Google Inc.
Available from: http://www.google.com, [Date Accessed: 6 August 2004]

Jung, J., Sit, E., Balakrishnan, H., Morris, R. (2002) *DNS Performance and the Effectiveness of Caching*, Massachusetts, Massachusetts Institute of Technology.
Available from: http://nms.lcs.mit.edu/papers/dns-ton2002.pdf, [Date Accessed: 31 August 2004]

Kaashoek, F., Morris, R., Dabek, F., Stoica, I., Brunskill, E., Karger, D., Cox, R., Muthitacharoen, A. (2004) *The Chord Project*, Massachusetts, Massachusetts Institute of Technology.
Available from: http://www.pdos.lcs.mit.edu/chord/, [Date Accessed: 31 August 2004]

Kerner, S. M. (2004) *With Release, X.Org Seals Fate of XFree86*, Connecticut, Jupitermedia Corporation.
Available from: http://www.internetnews.com/fina-news/article.php/3338031, [Date Accessed: 31 August 2004]

Krechmer, K. (2003) *Face the FACS: Proceedings of 3rd IEEE Conference on Standardization and Innovation in Information Technology held at Delft University of Technology.* Nederlands, Delft University of Technology. p.141-148.

Laoveerakul, S., Laongwaree, K., Tongsima, S. (2002) *Decentralized UDDI based on P2P*, Thailand, National Electronics and Computer Technology Center.
Available from: http://www.hpcc.nectec.or.th/C4/grid/UDDI.pdf, [Date Accessed: 27 August 2004]

Liu, X., Personal Communication 11 August 2004.

Lynch, J. (2003) *RSS News Readers Browse for You*, New York, Ziff Davis Publishing Holdings Inc.
Available from: http://www.pcmag.com/article2/0%2C4149%2C1267205%2C00.asp, [Date Accessed: 6 August 2004]

Maule, A., Personal Communication 13 August 2004.

Pau, R., Personal Communication 11 August 2004.

Schoechle, T. (2003) Digital Enclosure: the privatization of standards and standardization: *Proceedings of 3rd IEEE Conference on Standardization and Innovation in Information Technology held at Delft University of Technology.* Nederlands, Delft University of Technology. p.229-240.

Shafe, L. (2002) *The Shafe Report - Peer-to-Peer Computing*, Surrey, Software Opportunities Limited. Available from: http://home.btclick.com/lshafe/tsr2web.htm, [Date Accessed: 31 August 2004]

Smith, K., Personal Communication 11 August 2004.

Snowden, N., Personal Communication 11 August 2004.

Sun Microsystems (2004) *METEOR Distributed Hash Table*, California, Sun Microsystems Inc. Available from: http://meteor.jxta.org/, [Date Accessed: 31 August 2004]

Treloar, A. (1996) *Towards a user-centred categorisation of Internet access tools*, Victoria Australia, Deakin University. Available from: http://andrew.treloar.net/research/publications/NWS93/index.shtml, [Date Accessed: 6 August 2004]

UDDI.org (2002) *UDDI Technical White Paper* ,Massachusetts, OASIS. Available from: http://www.uddi.org/pubs/Iru_UDDI_Technical_White_Paper.pdf, [Date Accessed: 27 August 2004]

UDDI.org (2000) *UDDI Overview*, Massachusetts, OASIS. Available from: http://www.uddi.org/pubs/UDDI_Overview_Presentation.ppt, [Date Accessed: 27 August 2004]

YoOng, W.W.C., Personal Communication 11 August 2004.

Section 8: Critical Analysis of the Realm concept for widening access

8.1 Realm recapitulation with respect to the widening of device scope

It has been evident throughout this document that the Realm concept could be used to provide device context. This could provide a number of human and computational benefits because devices would be organised into logical groups with standardised interfaces.

Further research is required to determine if merely providing a "common framework" would provoke the widening of the Internet's device (*including information*) scope; although it was hypothesised, by the author. that the concept of context would provide important information and communication benefits which could provoke an increase in device scope.

To help summarise this, the example below illustrates three realms and their class schemas, see Table 8.1.

- The CCTV Realm:
  - is concerned with physical devices, whilst the "GeographicalData" and "CompanyData" realms are concerned with virtual devices.
  - extends the GeographicalData and CompanyData realms, which provide additional, optional functionality.
- The classes (*attributes, methods [including the codec types]*) merely provide an indicative structure for the purposes of illustration.

CCTV {Security-Public/Private}	***Realm Location (e.g. Camera1.Location.CCTV.Surveilance) – this address/location may be a unique identifier for the camera.***	
Extends: **GeographicalData; CompanyData.**		
Attributes	*Methods* (This is similar to a bean.)	*Supported Codecs (Inter-operable Plug-ins)*
Geographical Co-ordinates	.getGeographicalCoords(): int .setGeographicalCoords(int)	H.263 MJPEG
Proprietor (e.g. Directory.Company realm, uniqueID)	.getProprietor(): Realm, int .setProprietor(Realm, int)	
LiveFeed	.getLiveFeed(): Codec .stopStartLiveFeed(boolean)	
HorizontalOrientation	.getHorizontalOrientation(): int .setHorizontalOrientation(int)	
VerticalOrientation	.getVerticalOrientation(): int .setVerticalOrientation(int)	
GeographicalData {Security-Public}	***Realm Location (e.g. Portswood.Southampton.UK.CCTV.Surveilance) – this address/location may be a unique identifier for the GeographicData.***	
Attributes	*Methods*	*Supported Codecs (Interoperable Plugins)*

City/TownName	.getCity/TownName(): String .setCity/TownName(String)	HTTP XML – GeographicData Schema
Postcode	.getPostcode(): String .setPostcode(String)	
Latitude	.getLatitude(): int .setLatitude(int)	
Longitude	.getLatitude(): int .setLatitude(int)	
Altitude	.getAlitude(): int .setAlitude(int)	
CompanyData {Security-Public}	***Realm Location (e.g. Pseudocompany.Southampton.IT.Company.UK) – this address/location may be a unique identifier for the GeographicData.***	
Attributes	Methods	Supported Codecs (Interoperable Plugins)
Name	.getName(): String .setName(String)	HTTP XML – CompanyData Schema
Address Line1	.getAddress1(): String .setAddress1(String)	
Address Line2	.getAddress2(): String .setAddress2(String)	
City/Town	.getCity/Town(): String .setCity/Town(String)	
StateProvince	.getStateProvince(): String .setStateProvince(String)	
Country	.getCountry(): String .getCountry(String)	
InquiriesEmail	.getInquiriesEmail(): Email .setInquiriesEmail(): Email	
CompanySectors	.getCompanySectors(): String .setCompanySectors(String)	

Table 8.1

If a framework was created to facilitate the applications (*realms*) such as the above, then there might be an increase in the number of node devices. This is because it may only be possible to realise the benefits should a Realm contain a complete set of data entities (*nodes*).

However, although realms clearly provide a conceptual structure for physical (*including sensors*) and virtual devices, this type of idea isn't novel.

8.2 An alternative approach to Realms

IRISNET (2004) are developing a similar system which views all devices as sensors, which can interact with physical and virtual environments.

Gibbons *et al.* (2003, p.22) state (*with reference to IRISNET*) that one could view a PC's network interface as a rich sensor for the Internet virtual environment in addition to traditional off-the-shelf sensors such as webcams, microphones and motion detectors etc. They envisage a worldwide sensor web where users can query vast quantities of data from potentially millions of widely distributed, heterogeneous sensors. It is their view that the only aspect hindering access to such vast sources of information is a software architecture that realises this potential. Consequently, it is their aim to provide the missing software components.

With regard to widening access, Gibbons *et al.* discuss the fact that sensor network research, to date, has largely been focused on systems with severe resource constraints (*e.g. scarce energy, slow CPUs, limited communication speeds*) and where specialised hardware, operating systems, programming languages and databases have been designed. They state (p.2) that through IrisNet, they "*seek to broaden the scope of sensor networks to include wide-area sensor webs such as those comprising of widely-dispersed PC-class nodes with powerful CPUs that can process rich sensor data sources. A world-wide sensor web seamlessly integrates a wide range of sensor feeds, from high-bit-rate feeds from Webcam-equipped PCs to low-bit-rate feeds from traditional wireless sensor networks*." This, they state, will enable a variety of useful consumer-oriented services such as:

- Informing users when to approach the bus stop, or when water levels have become dangerously low,
- Waiting time monitors for reporting on queuing-delays at post offices, super-markets etc.
- Parking finders – finding a car parking space near a given destination,
- Lost-and-found services (*e.g. objects, pets and children*)
  – this could be used in health departments, defence, and computer monitoring services.

8.3 Areas to be accounted for when designing a sensor web and how they could be addressed by Realms

Gibbons *et al.* (2003, p.2) identify seven areas that need to be accounted for when designing such a "sensor web". Where aspects are implementation specific, JXTA has been used as the provider of the P2P network.

8.3.1 Planet-wide local data collection and storage

They state that sensors are necessarily distributed to observe the physical world. However, it is pointed out that because sensors would collect vast amounts of data, and because this data would typically be required to be retained for past and present observations, the system should store observations near the sources and transmit them across the Internet as needed.

In terms of realms, this requirement is adhered to quite well. Essentially a realm is a peer-to-peer structure based on human logic that specifies a node's attributes, methods, and codecs that it is required to adhere to in order to have membership of the ontology. However, in terms of node data, this is stored (*at least conceptually*) at the node (*peer*) itself.

It is envisaged that sensors with very limited hardware would be supported by sensor servers, which represent the sensors as if they had the necessary realm hardware requirements, e.g. a light sensor typically couldn't store past and present data locally, so a server would store the data and support the realm requirements – in terms of CPU cycles, bandwidth, memory etc.

8.3.2 Real-time adaptation of collection and processing

They state the sensor system should be able to reconfigure data-collection and filtering processes in reaction to sensed data. There are two distinct aspects to this:

1) It could be an internal arrangement (*i.e. part of the internal workings of the sensor*) and therefore is largely out of the scope of the Realms concept as it is focused within a specific type of sensor. For example, a water level monitoring sensor may be configured to sample once an hour with the exception of when the water level is above the emergency height when it senses every ten minutes.

Conversely,

2) Methods are an integral part of the Realm concept and merely the fact that a node supports them enables the potential reconfiguration by external systems. For example, a water sensor realm may enforce an increase sample rate method and any other collection/processing alteration methods. A sensor could continually inform an actuator of any change, which could control the sensor in real-time using methods.

8.3.3 Data as a single queriable unit

They state that the user should view the sensing device network as a single unit that supports a high-level query language and where queries operate over data collected from across the global sensor network. This is similar in principle to the way Google's search query encompasses millions of Web-pages.

This requirement has been one of the primary focal points for the justification of the Realms concept. Throughout this document it has been evident that realms limit the domain of interest, enabling the user and systems to view an ontology as a discrete system that supports a number of query features. At the most basic level, one could envisage an interface similar to that of Google, where the search is limited to a particular domain (*realm*). However, at a more advanced level, a realm could take advantage of XML node and data descriptions, which could be automatically organised and indexed in the storage aspect of a realm in a similar way to the Freenet storage system, see Freenet (2004).

8.3.4 Queries posed anywhere on the Internet

They state that users are "*accustomed to retrieving information stored anywhere on the Internet from anywhere on the Internet*" and that this ubiquity should be maintained. Further to this, the system "*should actively seek to exploit any locality between the querier and the queried data...because the sensed data are inherently coupled to a physical location*."

The Realms concept is firmly focused on node accessibility, where realm authorities are able to specify the security level (*assumed to be similar to Java classes: public, private and protected*), and nodes should be accessible by authorised users and systems from any Internet node. With regard to realms exploiting the locality between the querier and queried data, this ability is dependent on the peer-to-peer technology employed. It is envisaged that through a technology such as JXTA, a user or system would be able to directly connect to the node providing data.

8.3.5 Data integrity and privacy

They state that "*pervasive monitoring of the physical world raises significant data integrity and privacy concerns*" and that in their initial IrisNET prototype they assumed there to be a single universally trusted authority. They identify four constituencies with privacy concerns: service users, service authors, those in a sensed region, and those operating the sensor web.

The Realms concept is founded upon distributed authorities, where every realm has a governing group (*e.g. elected committee, company, individual, or overall peer consensus e.g. voting system*). With regard to data integrity and privacy, the requirement is largely outside the scope of Realms as this is primarily implementation specific. JXTA, for example, is cryptographic scheme neutral and supports many security mechanisms including some traditional methods. Transport Layer Security (TLS), the successor of SSL (Secure Sockets Layer), and X509.V3 digital certificates are just two examples of security mechanisms supported by JXTA, see Sun Microsystems (2002) and Yeager *et al.* (2002) for further details. Figure 8.2 provides the security advantages to be gained from P2P.

Intrinsic P2P Security Advantages

Advantage 1:	**Privacy** Since a message can be sent between two peers without going through a centralized server, there's no way an intruder on the server can read the message.
Advantage 2:	**No Central Point of Knowledge** Since content can be replicated indeterministically anywhere on a P2P network, it's impossible for an intruder to know the location of all copies. So content corruptions and denial-of-service attacks can't be performed.
Advantage 3:	**Web of Trust** When interacting with each other, peers can establish their own level of trust. They can evaluate the behaviour of other peers because there's no server in the middle to hide behind.
Advantage 4:	**Locality When Searching** For things, a peer will always ask another peer in its local domain first. As a result, bad behavior is limited to neighbours or direct contacts. When a server is contaminated, so are all its clients.

Figure 8.2

Source: Sun Microsystems (2004)

The above security mechanisms are focused on traditional Internet security which should satisfy service users, service authors, and those operating the sensor web (*or realm*). The final element is concerned with the people under observation within the sensed region. This area hasn't been studied in depth for this document as the Realms concept has largely been focused with node categorisation, location and interface standardisation. However, it has largely been assumed that the physical and virtual devices forming the realms have been implemented by the people and organisations concerned. For example, an IP phone or virtual film node located via the relevant realm would have been created by the proprietors, therefore there isn't any inherent privacy concern. However, monitoring sensors (*especially public*) are subject to privacy regulations which need to be adhered to. This is sensor specific and is not directly concerned with the Realm concept.

8.3.6 Robustness

They state that "*In a system that uses so many sensing devices and so many computing devices, failures will occur often. The system should operate smoothly despite these failures and the* *resulting unreachable hosts, unavailability of previously stored observations, missed new observations, and so on.*" Like many of the other factors described, this is essentially implementation specific. However, P2P networks are typically very dynamic and their nodes are often unreliable. Consequently, most P2P networks have stringent peer monitoring mechanisms to compensate and indeed Taylor (2003, p.5) states "*P2P systems need to treat failures as normal occurrences, not freak exceptions*".

JXTA is no exception as it provides both Peer Monitoring and Metering Peer Monitoring. Gong (2001, p.8.) states that these provide "*the capability to closely keep track of a (local or remote) peer's status, control the behavior of a peer, and to respond to actions on the part of a peer.*" He further notes that this capability is very useful should a peer network offer premium services as it can influence the reliability, scalability and guaranteed response time properties. For example, it may be better to shut down an erratic peer and transfer its responsibilities to another peer.

8.3.7 Ease of service authorship

The final aspect that Gibbons *et al.* identify is that the system should make it as simple as possible for service authors to develop services. They state that high-level abstraction mechanisms should hide the complexities of the underlying distributed data and query structures. This they identify as crucial to the adoption and proliferation of services using the system.

This requirement is largely out of the Realms scope as it is implementation dependent. However, at a high level, realms could ease authorship by providing realm standards, where nodes simply need to conform to a realm's specification. This means that the realm's interface etc. need only be written once, and nodes simply inherit the interface and conform to the attributes, method and codec standards.

This is also significant in terms of authorship of different systems, as a system requiring access to

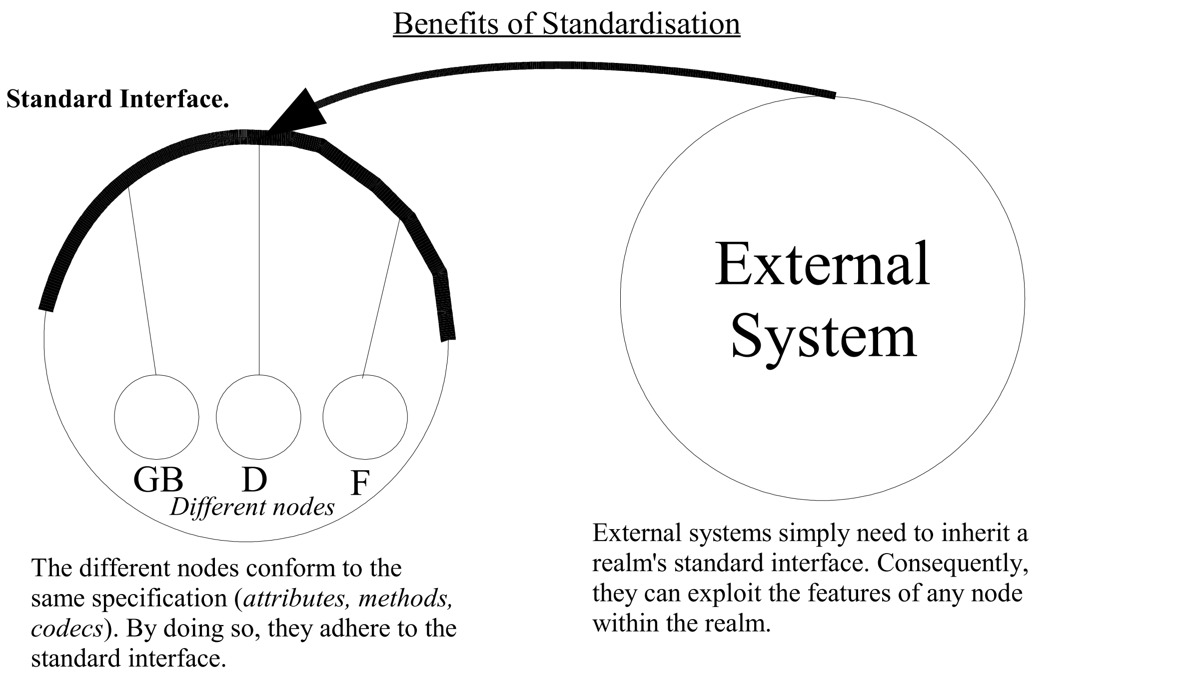

Figure 8.3

nodes within any given realm simply needs to conform to the realm's interface specification, where basic attributes and methods make all device data accessible to external systems in a similar way to current OO programming languages.

This eases the authorship of internal nodes and external systems.

An area where a similar type of plug-in infrastructure is emerging, is the Extensible Firmware Interface (EFI), where the PC BIOS is foreseen to be replaced by a modular, platform-independent architecture that can perform boot and other BIOS functions. Essentially it will be possible for hardware manufacturers to write a single EFI driver allowing any operating system (*that supports EFI*) to control the hardware through EFI functions.

Bibliography

Chen, R. and Yeager, W. (2002?) *Poblano, A Distributed Trust Model for Peer-to-Peer Networks*, California, Sun Microsystems Inc.
Available from: http://www.jxta.org/project/www/docs/trust.pdf, [Date Accessed: 19 August 2004]

Demeure, I. (2004) *An overview of P2P computing*, Cyprus, University of Cyprus.
Available from: http://grid.ucy.ac.cy/school/presentations/P2P-AnwireSchool.pdf, [Date Accessed: 21 August 2004]

Freenet (2004) *The Freenet Project - news – beginner*, California, Open Source Development Network Inc.
Available from: http://freenet.sourceforge.net/, [Date Accessed: 19 July 2004]

Gibbons, P.B., Karp, B., Ke, Y., Nath, S. (2003) IrisNet: an architecture for a worldwide sensor Web, *IEEE Pervasive Computing*, 2(4): 22-33.
Also available from: http://www.intel-iris.net/papers/pervasive-03.pdf, [Date Accessed: 17 August 2004]

Gong, L. (2001) Project JXTA: A Technology Overview, California, Sun Microsystems Inc.
Available from: http://www.jxta.org/project/www/docs/jxtaview_01nov02.pdf, [Date Accessed: 18 August 2004]

Hudson, P. (2003) What on Earth is EFI?, *Linux Format*, 40: 58-61.

Intel Corporation (2004a) *Intel Corporation - Extensible Firmware Interface (EFI) Home*, California, Intel Corporation.
Available from: http://www.intel.com/technology/efi/, [Date Accessed: 18 August 2004]

IRISNET (2004) *IrisNet: Internet-Scale Resource-Intensive Sensor Network Service*,
Available from: http://www.intel-iris.net/, [Date Accessed: 17 August 2004]

LEAD Technologies (2004) *Multimedia Converter and video codecs - Get your FREE Multimedia Converter*, North Carolina, LEAD Technologies Inc.
Available from: http://www.leadtools.com/utilities/video-codecs/codec_converter.htm, [Date Accessed: 17 August 2004]

Niemela, T. and Vega, L. R. (2004) *The Future of the BIOS: Modular, Scalable, Driver-Based Framework for EFI*, California, DeviceForge LLC.
Available from: http://www.deviceforge.com/articles/AT4903582708.html, [Date Accessed: 20 August 2004]

Sun Microsystems (2002) *Security and Project JXTA*, California, Sun Microsystems Inc.
Available from: http://www.jxta.org/project/www/docs/SecurityJXTA.PDF, [Date Accessed: 19 August 2004]

Sun Microsystems (2004) *Project Home Page*, California, Sun Microsystems Inc.
Available from: http://security.jxta.org/, [Date Accessed: 19 August 2004]

Taylor, I. J. (2003) *Peer Organization Lecture*, Cardiff, Cardiff University.
Available from: http://www.cs.cf.ac.uk/user/I.J.Taylor/DistributedSystems/PPTSlides/PeerOrganization.ppt, [Date Accessed: 21 August 2004]

Yeager, W., Williams, J. (2002) Secure peer-to-peer networking: the JXTA example, *IEEE IT Professional*, 4(2): 53-57.

Zimmer, V. (2004) *Advances in Platform Firmware Beyond BIOS and Across all Intel® Silicon*, California, Intel Corporation.
Available from: http://www.intel.com/update/contents/it01043.htm, [Date Accessed: 18 August 2004]

Project Summary and Conclusion

This project has studied

- The current and future device scope of the Internet
- The current and future Internet structure in both physical and virtual terms,
- The World Wide Web as an OO System,

It has also

- Proposed an object oriented concept, Realms, for the conceptual structure of the Internet.
- Questioned the opinion of potential end-users and experts on Realms.
- Studied realm interfaces, in terms of both conceptual location and physical implementation.

and it has criticised the concept of Realms for

- the Internet structure,
- widening access.

Section 1

The range of devices connected to the Internet is likely to increase as the benefits from a single communications means are realised. Two categories of device have been identified, physical and virtual. It has been evident from the reading that a number of devices will converge in terms of their abilities. The author of this document predicts that there will be a long term trend towards an increase in virtual devices, that is, intangible devices that can provide added functionality to the future convergent technology, e.g. assuming a PDA device, it could be possible to download software that is capable of extending the PDAs functionality for phone use. This is a trivial example, but the same principles apply to any system such as phone, mapping or virtual doctor services; functionality is increased.

Section 2

Largely to cope with the increase in devices employing Internet technology, the Internet support technology is planned to expand at a consistently high rate, see TeleGeography Research (2003). This includes new, faster and modular networking systems and the IPv6 virtual infrastructure that provides a vast number of uniquely identifiable addresses. It is largely believed that this will provide an infrastructure capable of accommodating vast numbers of sensors, helping people to learn about "how the world works".

From this study, Internet advancements may appear to be primarily focused on physical technology. This is not necessarily the case, as the World Wide Web Consortium and various academic institutions are actively working on a number of virtual Internet technologies, such as "agent technology". This is an incredibly vast subject and for a number of aspects, the author of this document is not sufficiently qualified to comment.

Section 3

From analysis of the likely future devices and infrastructure, the project focused on the conceptual needs of users and computer systems. Object Orientation (OO) was explored with a complementary view to other technologies. It was evident that many researchers thought that agent and the associative technology would play an ever increasing role within the Internet. However, the purpose of this research was to investigate the concept of Realms, founded on object oriented principles.

The closest work that was comparable to object orientation to conceptualise the Internet was Semantic Grid Ontologies. Pham *et al.* (2004, p.3) described the Semantic Grid as a means of enabling semantics to describe resources that computational systems could understand and potentially enable seamless automation. Their presentation highlighted the benefits of peer-to-peer (P2P) computing and collaboration.

Section 4

The concept of a Realm was defined and explained in terms of the current Internet infrastructure. The Realms abstraction system was perceived as a layer on top of the Domain Name System that could map logical nodes as opposed to merely physical nodes; that is a node was assumed to be any resource, physical or virtual. The conceptual categorisation mechanism was discussed and some of the OO principles were demonstrated for both physical and virtual devices. The potential benefits from this system were discussed including quality of service, security, the ability to promote and support abstract devices, and capacity to help facilitate drag and drop functionality to a device.

Section 5

Several interviews took place to try and determine whether there would be demand for such a system and if it was a feasible solution. Time was very much a limitation and consequently only five potential end users and two experts were consulted.

The end user reaction was primarily positive and all of those questioned stated that they grasped the concept reasonably well from the explanation given. Further to this, the users believed that the Realms concept was genuinely beneficial to them. From the two aspects discussed, search and standardised interface, the user reaction was mixed and implied that that both elements were equally important. The table below provides various realms suggested by the author and the average user desire for each.

Realm	***Average User Reaction (--No desire-- 1 2 3 4 5 --Very High desire--)***
Video	3.4
Book	3.8
Music	3
Shopping	4.2
Traditional Information Web-page	3.2
Telephony	3.2
Software	4.2
Abstract Family/Friend	3

It is evident that, collectively, there was desire for each of these realms; although it is clear from the raw results that there was a diverse reaction in some areas. A further aspect is that there might have been some bias in the results; the people questioned, although rating for a realm abstraction system, may have been influenced by the functionality aspect, an unfulfilled need, as opposed to the Realm structure supporting these facilities.

Further to the standard realms previously discussed, the potential realm users were asked if there were any additional realms that they desired. Two people had suggested a genealogy realm, one person suggested a medical Realm, and another suggested an interface to public services. These responses were quite surprising and it was clear that the end users felt that, although greatly useful, the current Internet structure, from their perspective, could be improved. One user even commented "*you typically have to search because the information could be anywhere...*" (West, *loc.cit.*) The user clearly thought that there were benefits to be gained from an ontology approach.

Section 6

It was evident throughout the document that the realm locations were important for discovery and access purposes. The Dewey Decimal System was compared with the Realm concept's requirements and it was evident that it would need to be altered to account for more diverse information sources because of the ability to supply often dynamic information automatically and nearly instantaneously from different environments. A final aspect was that license fees are applicable to the use of the Dewey Decimal System. This could make it difficult to use because of the financial element, especially when there are free alternatives. It was concluded that this area required further research.

JXTA was proposed as a likely platform for the implementation of the Realms concept. This was because it provided an infrastructure that abstracts the underlying P2P architecture, it is Open Source, it could support any type of device ranging from sensors to mainframes, and the system is built on open standards that don't favour any particular programming language.

Section 7

Current search engines, such as Google, Yahoo, Altavista and Lycos, offer categorisation services. However, it is evident that they are limited to web pages and are under private control which can stifle the development of this service. Further to this, a quick poll of six users found that out of the four that had used categorisation facilities, only one user found it to be a useful facility. This apparent lack of usefulness appeared to be evident in the Google directory service, which isn't immediately obvious from their index web page. It was evident from expert documentation that categorisation allows the user to explore a very limited portion of the large Internet information space. A further problem was raised in relation to semantics and categorisation. The vocabulary used throughout the Internet is diverse and categories can overload a user with information. However, from the research undertaken, it was evident that both categorisation and search mechanisms complement each other.

The Peer to Peer (P2P) architecture was largely assumed, throughout this document, to be the chosen architecture. This was because it expresses useful benefits, especially in relation to ontologies. However it was evident that such systems suffer from latency, lack of control and other problems. With regard to latency, P2P systems typically have large numbers of ‘hops’ between peers; that is, the number of nodes between peers is greater when compared to systems such as DNS. This is an intrinsic property because P2P systems form virtual networks that comprise of only a small percentage of the Internet computers. Further to this, information is distributed which can lead to further latency.

In terms of control, P2P systems are notoriously difficult to control because there isn’t a single owner and the networks are highly dynamic and the network is formed from a number of different sources which typically don’t have a permanent presence.

Universal Description, Discovery and Integration (UDDI) is one alternative to a P2P system for the implementation of the Realms concept. This is registrar focused in a similar way to current domain names. It supports, however, advanced descriptions enabling entities to list organisation, categorisation and services information. This could enable the formation of Realms whilst enabling governing authorities to maintain overall control of the network.

The exact choice is largely out of the scope of this document as it is Realm specific. There is no reason why centralised and decentralised architectures could not reside together; for example: Realm A could be centralised, Realm B decentralised and Realm C could be a hybrid of the two.

Section 8

There are two distinct aspects to widening access.

1. In terms of increasing the number and diversity of devices that connect to the Internet, it is unlikely that the Realms concept would increase either of these two elements. Many would argue that it is inevitable that the number and diversity of devices will increase regardless of the Realms concept. For example, it was discussed that the developers of IRISNET are working to help create a sensor web. However, although realms may not increase the long term diversity of devices, the time aspect is significant. If Realms or any similar concept proved to be popular, then it is likely that there would be an influx of devices with support for Internet connectivity.

2. The system could be used to widen access to information. Currently, information is distributed chaotically throughout the Internet resources. Search engines are used to mop up this information and supply appropriate information to user queries. If the Realms concept proved to be popular and became ubiquitous, system designers would design their systems with realm interfaces and would be more concerned with their system’s conceptual location. The combination of human and computer resources, and the open approach could provide ontologies and make information more accessible. It must be noted that if this approach to information organisation were to be successful, it is unlikely that this could be achieved without the semantic elements that the members of W3C are actively supplying.

Further to this, in terms of implementation and viability, further research is clearly needed, although JXTA and UDDI are significant systems that could help provide a foundation for the Realms concept.

Closing Passage

This document has served to give an insight into the potential benefits that the Realms system could provide. It was not the intention to provide firm answers or conclusions, but rather an attempt to provoke further thought into the conceptual structure of the Internet.

The ideas are far from revolutionary, and some suggest that such a system is simply not possible. However, if such a system did prove to be successful, it is likely that it would be because it could harness a “community approach” to solving the Internet organisation and conceptual structure aspects.

Further to this, the Realms concept could be used to promote competition and inter company co-operation. To help illustrate this significant aspect, read the following analogy:

Bed and Breakfasts (B&Bs), say within a town, are co-operative at one level, and hostile/ competitive at another. At a national level they may work together to promote the town/region, in terms of marketing etc. as it is in all of their interests that their sector is promoted. However, the B&Bs may be highly competitive at a micro level, where each B&B strives to entice customers through “their doors”.

This same principle could be applied to Realms. For example: a disk storage realm may contain numerous companies offering similar services, but all supporting the realm’s methods, attributes etc. The companies could work together to develop enhanced functionality and promote the realm (*sector*) in terms of advertisement etc. This could result in lower company operating costs, increased development for the companies concerned, and increased simplicity for the customer because there is a single sector to locate storage services. Companies could remain competitive by focusing on value added services, such as customer support and specialist applications of the particular realm’s technology (*e.g. a company specialises in mobile device storage capability, whilst another focuses on very high bandwidth services*).

End of document