

○ NETWORK AGNOSTICISM

WHY NETWORKS MUST DISAPPEAR

Belinda Barnet, Swinburne University of Technology

This article discusses pervasive computing as it relates to mobile content and mobile networks in particular, and argues that we need a unified networking concept that can bridge fixed and wireless networks and heterogeneous wireless technologies. It also discusses contemporary software applications that provide a content strategy that can transport a user's personal content and identity between networks, and calls for more work in this area. The article discusses current interconnectivity and interoperability between networks in Australia and recommends more work in this area.

'The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.'

(Weiser 1991)

In September 1991, Xerox Parc Chief Scientist Mark Weiser predicted that computing would disappear into the objects around us – into cars and coffee machines and refrigerators. These devices would all be connected to a network, and they would all have software systems that would implement 'ubiquitous applications' that can utilise the data we generate in our everyday lives¹ (Weiser 1991). He called this vision 'ubiquitous computing' and predicted it would emerge as the dominant mode of computer access over the next twenty years. An important part of his theory was that information processing would be thoroughly integrated into everyday activities, and that the objects we use or even carry (for example our credit cards or phones) would be network-aware and able to communicate freely with each other and with us. Since 1991, we have seen a proliferation of network-aware devices in our domestic lives, from Internet enabled refrigerators to in-car navigation systems and hand-held computers.

Many books and articles have been written about ubiquitous computing; it is now called pervasive computing, or more recently, 'everyware' (Greenfield 2006). Roberto Saracco and colleagues published a book (Saracco et al. 2000) that took Weiser's argument into the age of mobile technology and argued that telecommunications will soon become so integrated in our daily lives that we won't even notice it. Aside from future gazing, a good deal of money has also been spent on pioneering new devices. So has this vision arrived, and if it has not, what is standing in the way?

This article will argue that Saracco and Weiser's vision has not yet arrived, and that the main obstacle is not the lack of smart devices – it is the lack of a unified networking concept that can bridge fixed and wireless networks and heterogeneous wireless technologies, and the lack of a content strategy that can transport a user's personal content and identity between networks. This is an issue of interconnectivity and interoperability.

If we are to arrive at Weiser's vision of pervasive computing, then we need applications and content services that can freely adapt across multiple networks, technologies and devices. The device you are holding as you access your Facebook account should not care which network it is connected to – cellular or WiFi, WiMAX or 3G – the important thing is that you can easily access the service you want, anytime, anywhere and preferably on any terminal. Your home se-

curity system should be able to communicate with your mobile phone – even though one uses a cellular network and one uses a fixed line network. In both these cases, we have so-called ‘smart’ devices, devices that are network-aware: the goal is to enable them to communicate freely.

This vision of a world where people, content and also applications can cross freely between different types of networks is often dubbed ‘network agnosticism’ by technical development teams. Networks should be invisible as well as pervasive. They should be part of the background, like ambient music or the air we breathe. Applications and content should not care which network they are connected to or where they are when they do it – they must be able to adapt and move freely across heterogeneous technologies and even network generations.

This article will explore network agnosticism on both a technical level (i.e., the level of hardware, fixed line or mobile networks) and also content (i.e., the level of applications and systems). These two different levels pose substantially different technical problems, but both present challenges to the growth of pervasive computing. This article will be using the term ‘network agnosticism’ to embrace both levels because from a consumer’s perspective there is no distinction – it is all about the service experience.

In the contemporary era, although we are seeing a proliferation of ‘smart’ devices and technologies, the networks we use to access content are converging. Traditionally, users had mobile phones to make calls when they were away from home; landlines to make calls from home; broadband connections for surfing the Internet on a computer at work or home; televisions for watching broadcast information; radios for listening to newscasts and talk shows; and MP3 players for listening to music.

With the rise of Internet Protocol (IP) as the means to connect all these different devices to a network, we are seeing an increasing demand for interconnectivity between devices and applications. This is, of course, part of the old vision of ubiquitous computing Weiser outlined in 1991; one of the main requirements for the integration of information processing into everyday objects would be ‘a network that ties [all the devices] together, and software systems implementing ubiquitous applications’ (Weiser, 1991). In other words, true interconnectivity and interoperability in the devices we have in our lives. Pervasive computing implies the saturation of our environment with computing capability, yet ‘so gracefully integrated’ and context-aware that the interface itself becomes invisible.

This is obviously not our experience as users at the moment. Applications have trouble talking to each other and sharing content freely across networks. Devices have trouble sustaining different content services as they cross between network types (for example, fixed line or WiFi to cellular). Our home security system cannot yet communicate directly with our mobile devices when there is a problem, nor can we move freely as users between different networks and carry our identities with us. The various applications we use are not so ‘integrated’ in our daily lives that they can identify the content that might be useful for us at that time or in that place. We must open and login to various applications and networks several times a day. Our experience of computing is not one of seamless and invisible network connectivity – it is interrupted on many levels.

Let’s take social networking as our first example of the problem on an applications or ‘systems’ level. From a user’s perspective, social networks seem to be appearing (and disappearing) at a rapid rate – one minute it’s all about Friendster and MySpace, the next it’s Twitter, LinkedIn or Facebook. Each one of these new social networks requires a lot of time and effort: one must put up photo galleries, send invitations or requests to hundreds of friends, fill in all the boxes with

details about your life, install the various plugins and apps, and continuously feed the service with status updates.

Each social network requires several hours of work, not including daily ‘maintenance’ time. The amount of content generated (and in most cases duplicated) across different networks is staggering, particularly for the younger generation who communicate extensively on these networks. If only there were only a way of transporting an online or social ‘identity’ between networks, of re-using content you have created to support this identity. As keen social networker ‘krisis’ observes in a post on his long-running blog, *Crushing Krisis* (2009),

Friendster was the first prominent Social Network in America. Now it doesn’t even factor into the domestic conversation. People didn’t know that in 2003, so they gamely wrote their bios and uploaded their photos. Many of those people migrated to MySpace, where they posted more photos and wrote on a ton of walls. A lot of that same crowd also started to use FaceBook, where they posted yet more photos, wrote on a whole new network of walls, and penned pithy third-person status updates. For a single user the musical chairs of social networks can be mildly annoying. Do you even have your own copy of those photos? Do you really feel like hunting down all of those high school classmates again?

This is an example of a problem at the level of content services, or more precisely, social networks. At the moment, a user’s online identity (and all the content they have created to support that identity in the form of avatars, image galleries, status updates, music collections, and sometimes thousands upon thousands of friends or ‘fans’) cannot move freely between networks. In a truly network agnostic environment, your hard-won identity would be able to move with you between both applications and technologies. This of course raises its own issues around identity management and the privacy of personal data, as Satyanarayanan observes in his paper on challenges to pervasive computing: although the goal is make networks and applications invisible to users, one needs to ‘strike a balance between seamless system behavior and alerting the user to potential loss of privacy’ (Satyanarayanan 2001, 16). Users may not want their data to be shared between networks in all instances.

A US-based software company called mEgo are building a possible solution this problem. They have created a portable digital profile aggregator designed to integrate one’s online identity across multiple networks – an aggregator the user must consciously opt in to. This profile can be shared across different social networking sites – for example, Facebook, Twitter and Myspace. According to TechCrunch, the company recently finished a \$2.5 million funding round, and boasts that it has more than 1 million users. There are also a number of other portable identity aggregators in the pipeline – Gizmos and Voki, which both let users create portable talking avatars to post to their various networks (Kincaid 2009). mEgo is different from these other aggregators, however, in that it offers users the ability to store and share photos and profile information rather than just a fancy avatar.

This is important if we are to have content strategy that can transport a user’s personal identity between networks as well as just their picture or avatar.

The problem of interoperability also occurs at the level of devices and the networks these devices connect to, however. For example, if a user is watching a video on YouTube on their mobile device at home (connected to their own WiFi network), that stream will abruptly end when the user walks out the door. Why can't the device simply switch from the user's WiFi network to a cellular network without interrupting the content service?

To be a truly ubiquitous, 'ambient' experience, networks should be part of the furniture, they must exist in the background; they should enable our experience as users, not interrupt it. Although the two problems I have just outlined – moving between social networks and moving between, say, wireless and cellular networks – are obviously distinct on a technical level, they are very similar from a user's perspective. I wish to draw a parallel between them at the level of user experience.

Users would like their online experience to be seamless, and content they have created or wish to consume to be freely available and accessible on any terminal. Ideally, the service they are using would also be able to identify the user, and make useful content available to them based on where they are or what they are doing at the time. Needless to say, the ability to carry an online identity between networks would begin to facilitate this. Pervasive computing by definition implies that connectivity is available to users anywhere, anytime and on any device (Fan et al. 2004), and further, that 'ubiquitous applications' will be able to run on these devices and communicate with each other. Network agnostic technologies facilitate pervasive computing by allowing users to create or access content across any network, and to carry content with them across heterogeneous networks and technologies.

There is a good business case for creating network agnostic devices and services. A number of leading technology companies, including Alcatel-Lucent, Ericsson, BT and Nokia among others are involved in a European project called 'Ambient Networks'. The project is co-funded by the European Union (phase 1, which started in 2004, was granted 12.4 million euro by the EU) and it is industry-led (Cordis.europa 2008). The term 'ambient network' means a network integration design that solves some of the problems relating to switching between networks: it is by nature network agnostic, and in fact tries to establish cooperation between networks.

The EU Ambient Networks project began with the assumption that 'the future wireless world will be filled by a multitude of user devices, and wireless technologies', but at the same time, users will 'expect a rich set of communications services' and pervasive access to those services (Ambient Networks website 2009). Users will also want to transport content between networks. So the central problem is that on the one hand we have a proliferation of different technologies and networks, and on the other, there is a demand that they work together efficiently. The consortium was formed because the group felt this problem would not go away, and in fact it would get more pronounced going into the future.

To provide users with the services they want irrespective of their location, these networks have to cooperate. In a highly mobile environment, this network co-operation will have to be established 'on the fly'. The Ambient Networks project is about making this fast network cooperation a reality (Ambient Networks website 2009).

The consortium aims to do this by creating mobile technologies that provide access to any network (fixed, wireless, even satellite) through the instant establishment of inter-network agreements. They will establish this interoperation through a common control plane distributed across the individual, heterogeneous networks. This 'control plane' can be deployed as an integral component of new network architectures or as an "add-on" to existing networks (Ahlgren et al. 2005). It will be an interesting project to watch, but I think we should be considering projects of this kind in Australia – not just to create new technologies, but to change the way we think about and design content services.

Part of this change requires us to think about what kind of telecommunications future we would like to see in Australia. If this future involves ubiquitous or pervasive computing, then we also need to think about what is blocking our progress towards this vision. After the EU Ambient Networks project, I think it is reasonable to assume that we will see a further proliferation of smart devices and wireless technologies in Australia, but at the same time, we will see a strong demand for interoperability between these devices and the networks they are connected to; if we can make any assumptions about future technologies, this would have to be it. We will see cheap computing "spread around like grass seeds", and connectivity integrated into everyday devices (Greenfield 2006, 11–12). This being the case, it makes sense to design and think about technologies that will enhance network cooperation and bridge fixed and wireless networks and heterogeneous wireless technologies. This is a world that will be saturated by computing and network-capable devices, yet we will expect our devices to know us better and to be 'aware' of services and content that might be useful for us at that time. We can start towards this future by demanding that new systems, new applications and new devices are network-agnostic.

ENDNOTES

- ¹ This is sometimes called context-awareness, meaning a service that understands the context the user is currently in (eg their location, what they have previously been doing, etc). I do not have room in this article to explore this further.

REFERENCES

- Ahlgren, B; Eggert, L; Ohlman, B; Schieder, A. 2005. '*Ambient Networks: Bridging Heterogenous Network Domains*', Ambient Networks Project.
- Ambient Networks website. 2009. Accessed August 2009. Available from: <http://www.ambient-networks.org>.
- Cordis.europa. 2008. 'WWI Ambient Networks Project Fact Sheet', http://cordis.europa.eu/fetch?CALLER=FP6_PROJ&ACTION=D&RCN=71236&DOC=8&CAT=PROJ&QUERY=1.
- Fan, L; Akhtar, N; Chew, KA; Moessner, K; Tafazolli, R. 2004. '*Network Composition: A Step towards Pervasive Computing*', Proceedings of the 5th International Conference on 3G Mobile Communication Technologies, London, UK.
- Greenfield, Adam. 2006. *Everyware: the dawning age of ubiquitous computing*. New Riders. pp. p.11–12.
- Kincaid, J. 2009. 'mEgo Raises Another \$2.5 Million for its Multimedia Social Avatars', January 21, TechCrunch. Available from: <http://www.techcrunch.com/2009/01/21/mego-raises-another-25-million-for-its-multimedia-social-avatars>.
- Saracco, Robert; Harrow, J; Weihmayer R. 2000. *The Disappearance of Telecommunications*. Wiley, IEEE Press.

Satyanarayanan, M. 2001. 'Pervasive computing: vision and challenges'. *IEEE Personal Communications* 8(4): 10–17.

Weiser, M. 1991. 'The Computer for the 21st Century' in *Scientific American*, February.

Cite this article as: Belinda Barnet. 2009. 'Network agnosticism: Why networks must disappear'. *Telecommunications Journal of Australia*. 59 (3): pp. 40.1 to 40.6. DOI: 10.2104/tja09040.