Mobile Data, mBusiness and Machiavelli

The nervousness in the investment community concerning the future health of the 3G license winners is being compounded by reports that Wireless Local Area Network, or WLAN, solutions could steal their market. In this paper we use historical evidence and practical experience to get a balanced view of the real issues and to present a less pessimistic prediction of the future for the 3G licensees.

One of the common features in all of the science fiction epics to have graced the silver screen is the mobile communication device that allows our hero or heroine to get vital information. This device is much more than just a mobile phone – it is the portable part of a sophisticated network that connects the roaming user to the resources they need.

The technology that can make this piece of science fiction a reality is already with us. It has yet to be fully developed and deployed but all of the key components exist; the data capability of mobile networks is evolving quickly and the Internet is established as a worldwide information network. In this paper we consider the evolution of the mobile Internet and suggest that two key forces, neither of them altogether obvious, will shape it.

The first of these forces is connectivity. It may seem rather mundane, but history tells us that it will be the carriers, not the content providers who determine the shape of the mobile Internet. First and foremost, people want to communicate. Access to information is a secondary requirement.

The second force is pragmatism. For all the investment in third generation mobile systems as an all-embracing solution, it seems likely that the mobile Internet will be a mix of technologies. In this paper, we paint a picture of a hybrid network that combines centrally managed 3rd generation mobile networks with distributed LAN technology. The balance between the two will be determined largely by economics – for any given application the cheaper bandwidth will drive out the more expensive.

Mobile Data

In the early 1980's, the BBC introduced a series called 'Space Cops' that rapidly built a cult following. One of the most memorable features of the series was the device that accompanied the lead character everywhere he went, his personal Oracle. This interactive unit provided all the information, sent all the messages and carried out all the analysis that any self-respecting space cop could wish for. Not only did it provide an essential link into the network, it also told the network where the Oracle (and, hopefully, its user) was located.

It is important not to confuse this sort of mobile computing and networking with the portable computing and networking that has been possible on earth for some years. In this paper we are looking at a world where activities are not disrupted when the user moves from one point of attachment to another. Instead, all of the necessary reconnections occur automatically, behind the scenes, as an inherent feature of the system.

The obvious question to ask at this point is 'so what?' Why should anyone be at all interested in doing something on the move that is more readily achieved from a fixed location. The simple answer is that, given a choice, you wouldn't. Ideally you would want to sit down and read your e-mail or download a file at a fixed location. But sometimes you don't have that choice. There will be occasions when you are travelling and still need access to information and, therefore, need to establish a roaming connection to a network.

So lets go on to consider what sort of mobile network links are really needed.

Content is King?

Over the last few years, the wireless industry has done very well selling low bandwidth pipes to users. However it now appears to be about to repeat a classic mistake of previous communications technologies by concentrating on content, rather than connectivity. The third generation (3G) systems being planned for introduction around the world over the next few years promise considerably higher bandwidth than current offerings. This bandwidth is universally touted as a way to provide Internet access, and in particular to sell content (that is, access to a fixed source of information) to users.

Yet the Wireless Application Protocol (WAP), which was designed to deliver content to wireless devices, has been a huge disappointment. By way of contrast, the Short Message System (SMS), which offers low bandwidth digital messaging between users, has surprised observers by its success. For example, in the U.K., between the second quarter of calendar year 1999 and the second quarter of 2000, the number of SMS messages grew from 159 million to 1.42 billion [source: Oftel].

This should come as no surprise to any student of Machiavelli, who devised all of his strategy based on what people actually do and what really happens, rather than what seems attractive. Throughout the history of communications, connections between people have consistently been more important (and lucrative) than access to information. For instance, the U.S. postal system of the 1800's was awash with Newspapers that needed to be delivered. This was the dominant content, at least in terms of volume, weighing about 20 times as much as the letters carried. But it was the letters that brought in most of the money needed to run the postal system – around 85% [ref: John]. On the Internet, it is still email that is king, even if its volume is relatively small.

In both instances, the content may be dominant in terms of volume, but the real demand, and hence most of the money, is in the point-to-point communication.

Looking at the current world of mobile communications, SMS provides connectivity and WAP provides content (or, more accurately, a means of accessing information). Following the Machiavellian principle that people's reactions in given situations are consistent it should have been completely predictable that SMS would prove more popular than WAP. Things that think want to link!

What does this mean for the cellular operators and their suppliers? In broad-brush terms, they should be striving to provide the connectivity people need, rather than the content or 'killer application' that they might want (or think they want). In other words, they should stick to what they are good at! Mobile Business (mBusiness) may capture the headlines and the imagination, but the operator who can offer the underlying network is likely to capture the money. The next question, of course, is what sort of network does an operator need to invest in?

Rabbit's Revenge?

The popular contender to deliver the sort of connectivity that supports high speed Mobile Data (and, by implication, provides the platform for m-Business) is 3G. A huge amount of resource has been and continues to be poured into the provision of third generation services. As one would expect with a mainstream telecommunications technology, much of this resource is focussed on agreeing standards, making sure the network is scaleable and attending to operational and support issues. 3G, when it arrives, will be of 'industrial strength' [ref: Holma]

But the prospect of 3G cellular providing the universal high-speed data service is being widely questioned [source: Woolcock] and alternatives, notably Wireless Local Area Networks, are being suggested. In this section, we examine the characteristics of both technologies and suggest how they might complement one another. The table below summarizes the basic capabilities of technologies that seem the most likely contenders to provide the mobile Internet.

Wireless Technology ranges when deployed in Urban areas, providing indoor and outdoor coverage.		
	Max Speed	**Range**
3G Urban Cell	384 kbps	500m
Bluetooth 1.1	1Mbps	<10m
Bluetooth 2.0	2Mbps	50m
IEEE 802.11b	11Mbps	100+m

The focus in the remainder of this paper is on 3G and IEEE 802.11b, as they provide network connectivity, as opposed to Bluetooth, which is, essentially, a wireless cable substitute [ref: Angell].

The technical challenge of building a next-generation cellular network is proving difficult for even the most advanced wireless carriers. Examples of this include NTT DoCoMo's suspension of their FOMA (Freedom of Multimedia Access) launch at the last moment and British Telecom's recent delay of the Isle of Man test bed until later in the year. These setbacks, combined with the huge cost burden of the 3G licenses, have allowed other technologies to put the squeeze on 3G. In particular, enhanced data capability on existing mobile networks has the advantage of lower deployment costs and Wireless LANs are being designed to interconnect portable high-speed devices over limited (but not trivial) distances.

In many ways the existing GSM/GPRS network promises to provide what many users want, which is good voice service, a modem speed data connection and SMS. Wireless LANs, with their high speed (around 10Gbps) and proposed metropolitan coverage in some markets (up to 40km radius with multiple base stations, each with a practical outdoor range of around 100 metres) will satisfy another large part of the user population. In particular, they appeal to those who want a fast connection to local services, (which is what typifies mBusiness).

And so, in truth, 3G on its own is probably not going to deliver the panacea that many had hoped for. A number of reasons have been cited for this – shared transmission media reduces speeds from theoretical Mbps to a more realistic few hundred kbps and battery technology will set a fairly low limit on the time that a mobile phone can be active (and will probably cause it to be rather too hot for comfort).

But these limitations will constrain any solution. So, if there is not going to be a panacea for a New World of mobile applications, what is likely to happen?

Taking history as our guide, the probable way forward will be for several strands of technology (some existing, some new) to come together, and for market needs to evolve to take advantage of the developments possible under such a scenario.

Looking back through time, it is always easy to see why particular developments happened when they did, as alternative scenarios are neatly eliminated from consideration by the course of history. However, looking to the future, it is very difficult to predict what will succeed. There are many credible options, and it generally takes two or three random developments (often in unconnected fields) to come together in a fortuitous way to make one option take off, another crash.

Perhaps the best way to see the road ahead for mobile data is to look again at the need – connection not content. Without a simple, inexpensive connection, acceptable to the waiting populous, there will be few services on the emerging market. When it comes to connection, there are two sure-fire winners in the data market – Ethernets and IP networks.

And so to the unique feature of mobile business mentioned (fleetingly) already – it is largely dependent on and driven by user location. This means that it is, predominantly, a localised service and that rapid fulfilment is key – a roaming user is not going to stay in one place long enough to wait for a centrally dispatched FedEx van to deliver what they want. This means that a significant amount of communication is likely to be short range, between the user and the nearest provider. Of course, the full range of mBusiness communication via the Internet is also necessary and so the network must cover the wide area. But the local area is the fundamental building block.

Wireless LAN technology (notably the IEEE 802.11b standard which is probably ahead of its rivals, such as HiperLAN) is a current favourite to meet the demands of mBusiness. But is this just the next example of technology hype or is there a real business case for it? The bulk of press coverage promoting the 802.11 standard has tended to concentrate on the technology itself and, while many reports quote headline grabbing potential revenue figures, there is rarely any attempt to outline how this income might actually be realised.

We can gain some practical insight here by looking at Finland's Jippii subsidiary Wireless Network Solutions as they are aiming to provide wide area coverage using IEEE 802.11 based technology. In fact, Jippii already operates a high speed Internet access solution, called

Freedom, in more than 100 locations in 50 Finnish cities with plans to rollout the service overseas. Service is currently offered in business districts, business centres, hotels, airports, restaurants and multi-tenant office buildings. The first extensive urban area has been covered in Seinäjoki, with plans to cover other residential areas such as south of Espoo.

If this type of solution is going to be a success anywhere, then Finland (with its seemingly insatiable appetite for all things mobile) is probably a fairly safe bet. While accounts of Jippii's Freedom solution are strong on technical capabilities, there is no mention of the number of users it has attracted or, quite crucially, how they are billed.

The key issue here is that standards such as 802.11 were developed as an alternative to traditional cabled LANs, notably the Ethernet, and only define the air-interface between a wireless modem and a base station hub. When used in the environment for which it was designed, i.e. private LANs with known users, the solution is perfectly adequate, but introduce public users and a whole raft of new issues emerge. The diagram below illustrates some of the support facilities that would need to be provided.

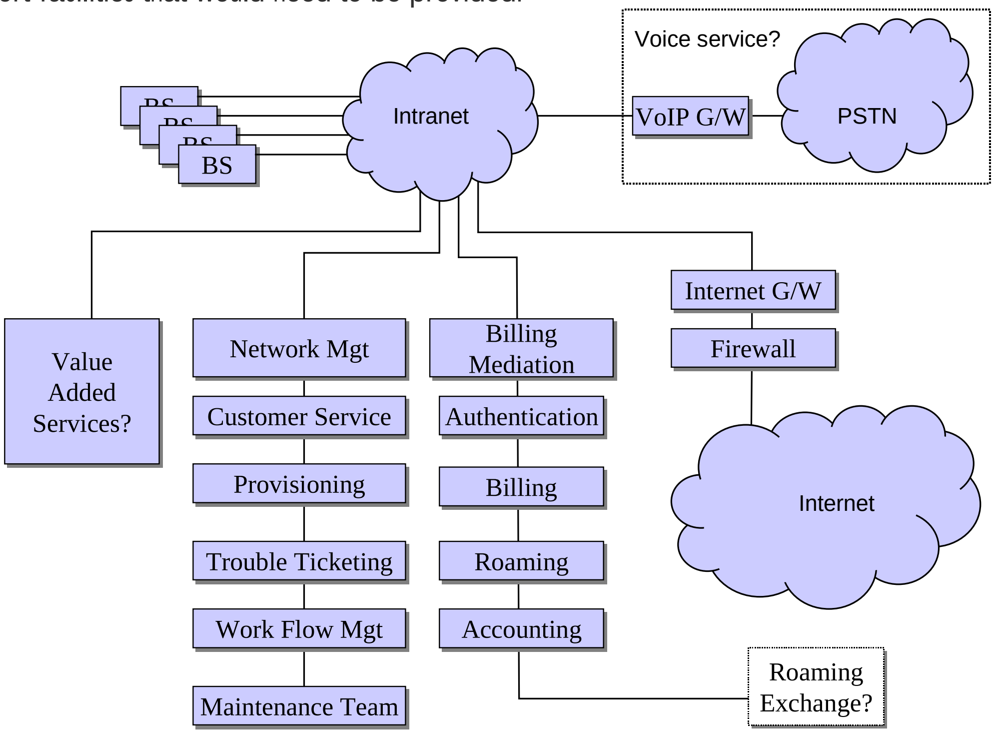

In the diagram, the users attached to the base station (top left) are part of a public access system. For any sort of commercial service, they must be provisioned, managed on an ongoing basis and, most importantly, billed. The facilities to carry out these tasks are part-and-parcel of cellular networks, but proprietary implementations would need to be developed for WLAN solutions. Articles comparing implementation costs often (conveniently) overlook the cost of this. The importance and complexity of operational support systems should not be underestimated [source: TOM].

So, for all of their appeal, Wireless LANs are unlikely to supplant cellular networks in the short term. They will doubtless provide part of the mobile Internet, but there are practical issues to

be addressed before viable end-to-end services can be offered. And a lot of design effort needs to be expended along the way.

The New Order

Given our brief assessment of candidate technology, let's start to paint a picture of everyday life in the near future. Most people will be carrying around a handheld device that is triple homed so that it can talk 3G to a mobile operator, IEEE 802.11b for a local network provider and Bluetooth as a cable replacement.

High-speed local connections will be routinely used for visual communications and other content-rich applications. Remote connections would be provided through an Internet gateway on the Wireless LAN. Since these usually require you to sit down the connections will be made inside the home or office, or near to some (well-appointed) conurbation – all of which is well suited to the use of a wireless LAN. This may be a slightly restrictive view but just imagine someone playing multi-user Quake on their Gameboy whilst driving. On second thoughts, please don't!

When people move around, they would probably have to use a wide area connection, linking to one of the core 3G (or, perhaps, GPRS) networks. This is good enough to pick up a specific (and usually temporary) piece of information that you cannot wait for. Some data, like travel alternatives or weather conditions at an intended destination, are needed 'on the move'. But most of the time the choice would be to use high-speed LAN access, which provides the necessary bandwidth (and carries the 80% of traffic which is local).

How do these fit together? The full answer to this is yet to emerge as many of the operational issues raised in the previous section (e.g. end-to-end billing, fulfilment, assurance) have yet to be resolved. At the networking level both technologies peer at the IP level, which can give a specific user seamless access to both central and remote services, as shown below.

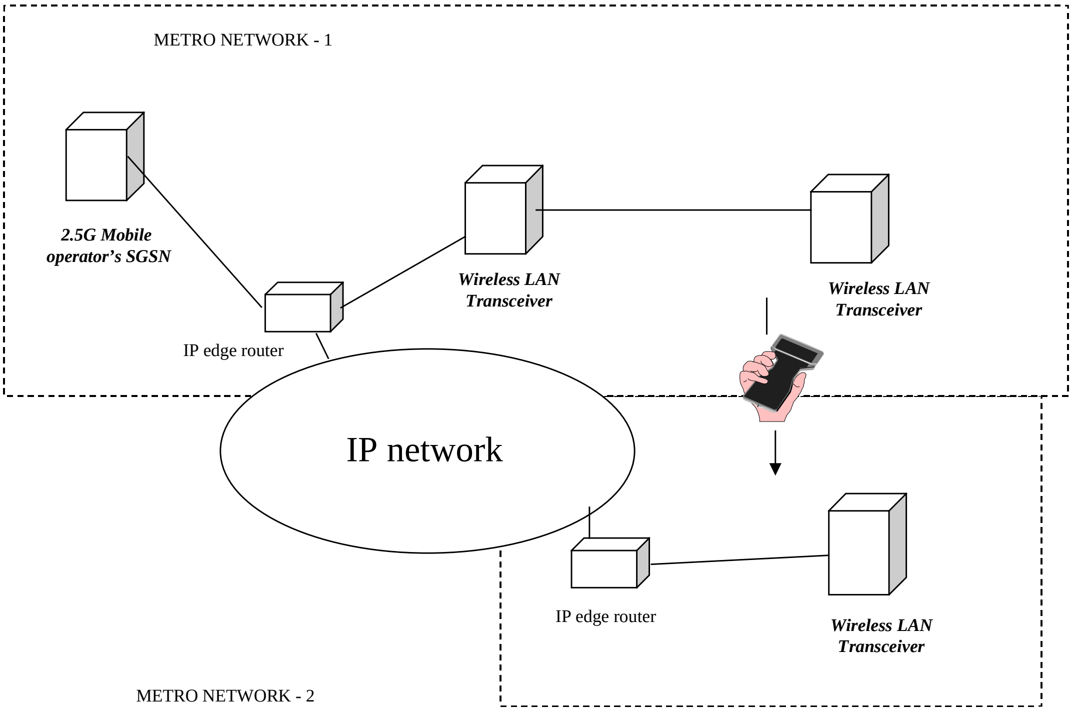

In this illustration, it is the Mobile IP protocol that provides the ‘glue’ that manages the end to end connection as the user roams from one place to another. This is discussed in some detail in the panel overleaf.

Reality Bytes

In theory, theory and practice are just the same. In practice, they’re not. There are many practical issues in the deployment of any network. This section considers a few of the areas where unwary providers could easily find themselves mugged by reality.

Comparing technologies often helps to highlight the potential implementation costs of providing the network coverage needed to support mBusiness. The high bandwidth that IEEE 802.11 offers and the requirement to reduce the risk of interference in an unlicensed spectrum band will realistically limit the coverage from a base station to a few hundred metres. This is a comparable range to technologies such as DECT and CT2 in Europe and Japan's PHS (Personal Handyphone System). The CT2 system failed in the UK because the technology and service had already been superseded by cellular, but PHS did enjoy a brief period of success in the Japanese market. However, in order to provide coverage to Greater Tokyo, 90,000 base stations were required. This historical perspective gives some indication of the sort of numbers needed to provide the city-wide coverage needed to host useful services.

Is this sort of coverage practical? Keith Woolcock, in his report “Barbarians at the gate—Wireless LAN storms 3G citadel,” states that the “DM 97 billion raised in the German 3G auction could have bought 60m Wireless LAN base stations”. So the economics are not out of line with real life….but there is more to providing a service than simply installing enough connection points.

Because both local and wide area connectivity are required, the issue of backhaul provision must also be addressed. So each base station site must be provided with a link to a backbone network, which ultimately supports the connectivity to the Internet. In order to provide a reasonable grade of service (commensurate with that planned for 3G services) each connection must be of the order of 2Mbit/s, bearing in mind that this capacity must be shared between all simultaneous users of that base site. This all adds to the cost of deployment. For example, BT charges £2,200 pa for a link of this capacity within central London. If the operator wishes to take full advantage of the much vaunted 11Mbit/s capability of the IEEE 802.11 standard, then a much larger backhaul pipe will be required – one or more 34Mbit/s connections, which currently cost around £30,000 pa from BT in central London. If it is assumed that a city like Greater London can be covered with half the number of sites of Greater Tokyo, then the annual cost of backhaul alone will be £100m for 2Mbit/s rising to £1.4bn for 34Mbit/s. And this only provides coverage in a single city. There will undoubtedly be cheaper options for deploying a full service solution, but it is clear that deployment costs – often overlooked – are significant.

In addition to the creeping costs of deploying a solution, there are other issues that potential operators must consider. For instance, the frequency band in which IEEE 802.11 operates is unlicensed. This means that there are no licence fees to recoup but it also implies that the air interface is open to multiple uses, including appliances such as microwave ovens and, critically, Bluetooth to which the IEE802.11 experts are currently trying to develop resilience. Will operators risk offering (and indeed, will people be prepared to pay for) a commercial service over which they ultimately have no control of the quality of service?

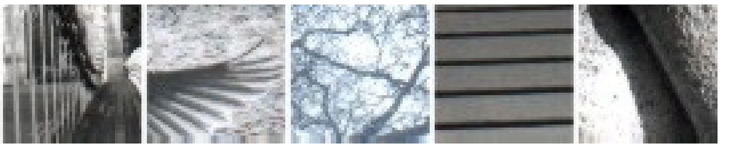

The key to the Mobile IP protocol is that it uses two addresses to ensure that a network connection is not broken as its user roams from one place to another. The first address is the user's "static" address, in essence his "home" address. The second address, the "care-of" address, changes with the point of attachment to the network.

Since every mobile device has its own unique IP address, if the user has moved away from his home network, the home network forwards the packet from the home address to the care-of address and to the mobile device. The mobile device uses a "Home Agent," which is a specially designated server that takes responsibility for intercepting and forwarding packets for absent subscribers. It registers with the home agent through the use of a "Foreign Agent" which is in the foreign network. The Foreign Agent's job is to care for visitors on its network by finding a visitor's home network and informing it that it has temporary care of that visitor.

As you might expect, the Mobile Internet Protocol is a bit more complicated than this simple explanation. For example, the way it routes packets from the Home Agent to the Foreign Agent is through encapsulation. Why? Well if the routers were aware of the "native" packet, they would see the Home Agent address and send the packet back to the Home Agent. So, the packets are encapsulated with the address of the Foreign Agent. When received by the Foreign Agent, the temporary "wrapping" is removed to reveal the original address that is then sent to the "roaming" recipient. Another example of the complexity of Mobile IP is that it uses "cache agents" to deal with the fact that packets must first be delivered to the home network which can make the end-to-end routing inefficient.

Anyone familiar with cellular networks would, no doubt, see a lot of similarity between the Mobile Internet Protocol and roaming between cellular networks. The following diagram explains exactly how a packet gets to a mobile device that is visiting a foreign network:

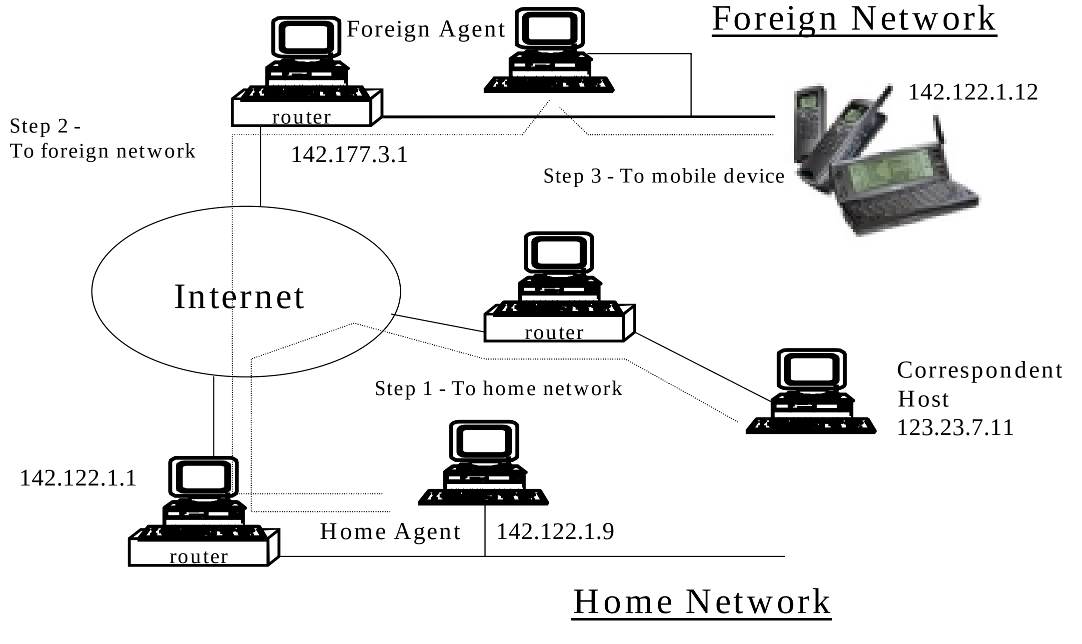

Here's how it works. The mobile device, IP address 142.122.1.12, gets a message from a computer, called a "correspondent host," on another network. The IP headers of the packets leaving the correspondent host have the same IP address. When the packets reach the home network, the home agent inserts an additional IP header and forwards it back into the network. The new address is 142.177.3.1, which routes it to the foreign network being visited. In the foreign network, the registration process should have informed the foreign agent about the presence of the mobile device so that when the encapsulated packet arrives, it knows to remove the outer header to reveal the original IP address and then forward the packet to the mobile device. Of course, return messages from the mobile device to the correspondent host proceed as normal messages.

An issue that has been skated over so far is that of ownership – who provides which part of the overall network. If the 3G service provider releases enhanced mobile devices and PDAs, users will expect these portable devices to work in the wide area in the same manner as today's crop of phones. In this case, the Wireless LAN access point is to 3G what DECT/CT2 (launched as Fido in Italy, Bibop in France and Rabbit in the UK) was to 2G mobiles and therefore risks the same fate. In other words, there is a prospect for a dominant player to own the user and thereby suppress competition. Of course, if wireless modem cards are GPRS/WLAN (and after 3G launch, GPRS/3G/WLAN) this enables the portable device to switch to the highest speed (or cheapest) option at any location and allows the most suitable operator to be selected.

Finally, the difficulty of providing operational support should not be underestimated. This would be particularly true in a hybrid network, where customers expect services (and associated charges) to be end to end, not partitioned on the basis of underlying technology. A possible solution here would be for customer care and billing to be handled through the standard 2G/3G mechanisms, with the Wireless LAN coverage appearing as just another network cell. This has the advantage over the independent operator model, as that requires a 3rd party (i.e. the local operator) to bill for local access on top of any service charge/subscription the user is already paying.

Conclusion

The contention in this paper is that raw connection capability, not content, will be the lucrative part of the mobile data business. Furthermore, Wireless LAN and 2.5G mobile technology as well as 3G will provide this connection capability.

The deployment of this hybrid mobile data network will provide the conditions necessary to make mBusiness viable and will promote its development. If basic cheap high-speed connectivity is available then history shows that there will be lots of people who use it to offer new services. A key point for the operators is that they should not try to control all the value chain and thereby suppress initiatives as they would already have an excellent position from which to gain a significant share of a potentially large market.

Our conclusion is that both individuals and businesses will opt for the most economic way of delivering an adequate solution; pragmatism will drive progress, not ideology and there will be a mix of solutions on offer. If this means different technology for local and long distance traffic or separate handheld devices and mobile phones, so be it. A large part of the paper addresses the practical considerations and issues that need to be addressed. We have asked some of the key questions and suggested where answers might lie. Real progress will be determined through compromise, barter and good design.

As Ira Brodsky, president of Datacomm Research observes, "public wireless LAN operators must join forces with 3G mobile phone carriers to achieve necessary coverage and service bundling. Likewise, third generation mobile phone operators need public WLANs to offload heavy indoor traffic from their lower speed, wide area networks".

mark.norris@intercai.co.uk **Tel: +44(0) 1628 478 470**
adrian.golds@intercai.co.uk

Intercai Mondiale Ltd, Regatta House, High Street, Marlow, Bucks SL7 1AB, UK
www.intercai.co.uk

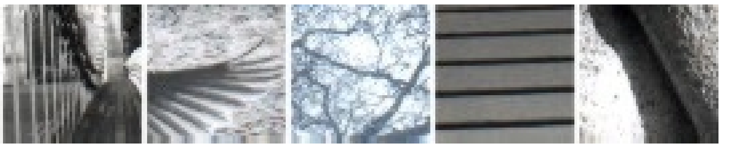

References

Anthony Jay, *Management and Machiavalli; power and authority in business life,* Hutchinson Business Books, 1987

Harri Holma & Antii Toskala *WCDMA for UMTS; radio access for third generation mobile communications*, John Wiley & Sons, 2000

Telecommunication Management Forum, TOM – *the Telecommunications Operations Manual v2.1*, www.tmforum.org

R. R. John, *Spreading the News: The American Postal System from Franklin to Morse*, Harvard Univ. Press, 1995

Jack McCue, *The Mobile Internet Manifesto: 4G Networks of the World, Unite!*
Economist Conference on The Wireless Internet, Stockholm, Sweden, May 2001

Mark Norris, *Mobile IP technology for M-Business*, Artech House, 2001

Andrew Odlyzko, *The history of communications and its implications for the Internet* available at http://www.research.att.com/amo

Oftel, the U. K. Office of Telecommunications, Nov. 2000 *Market Information* available at http://www.oftel.gov.uk/market/miu1100.pdf

Keith Woolcock et al, *Barbarians at the Gate – wireless LAN storms 3G citadel*, Nomura Equity Research, 2001

Chris Angell, *Bluetooth – the 3G enabler*, Intercai Mondiale White Paper at www.intercai.co.uk