

March 2015  
Geoff Huston

## The Mobile Internet

It has been observed that the most profound technologies are those that disappear (Mark Weiser, 1991). They weave themselves into the fabric of everyday life until they are indistinguishable from it, and are notable only by their absence.



The feat of reticulating clean potable water into every house, so that it is constantly accessible at the turn of a tap is a great example of the outcome of large scale civil engineering projects, combining with metallurgy, hydrology, chemistry and physics. But we never notice it until it is no longer there. Similarly, the adoption of household items such as the refrigerator, washing machine and stove. And of course let's not forget the feat of the domestic electricity supply grid. Prior to the construction of national grid systems, electricity's customers were entire communities and the service role was in lighting the town's public spaces at night. Domestic electricity was an unaffordable luxury for most households. Today we simply take it for granted.

Computers are also disappearing. Today's car has more than 100 million lines of code running in more than 100 microprocessor control modules. What we see in the car is not these devices, but instead we have cars with anti-lock brakes, traction control, cruise control, automatic wipers and seat belt alarms. Each of the car's underlying control systems are essentially invisible, and about all we get to see is the car's human interface system. This visible system, essentially an entertainment controller and navigation service, is currently the space where both Apple and Google are jostling for position with the auto makers, while all the other microprocessor systems in the car remain unremarked and little noticed.


So how should we regard the Internet? Is it like large scale electricity power generators: a technology feat that is quickly taken for granted and largely ignored? Are we increasingly seeing the Internet in terms of the applications and services that sit upon it and just ignoring how the underlying systems are constructed?

What about the most recent Internet revolution, the massive rise of the mobile "smart" phone? Will the use of a personal mobile computing device be a long lasting artefact, or will it be superseded in turn by a myriad of ever smaller and ever more embedded devices?

What should we make of the mobile smart phone industry? Is this all-in-one device headed down the same path of future technology obsolescence as the mainframe computer, the laptop and even the browser?

Or are these devices going to be here to stay?

## How did we get here?



One way to answer this question is to look at the evolution of the computer itself. Computing is a very new industry. Sure, there was Babbage's analytical engine in the 19<sup>th</sup> century, but the first computers appeared in the mid 20<sup>th</sup> century as programmable numerical calculators. These were massive feats of electrical engineering, built at a cost that only nation states could afford, and were of a size and fragility that they required their own building, power and conditioned environment. Valves are large, require high levels of power and are fragile. For the early computers, such as ENIAC, a small cadre of folk

were employed to keep these behemoths running. This model of computing was one that was only accessible for a few, and at a cost that was completely unaffordable for most.



The invention of the transistor changed everything. Transistors were far more robust, used far less power and could be produced at far lower cost, leading to the advent of the commercial computer in the 1960. These units, such as the ubiquitous IBM System 360, were used in large corporates and in universities and research institutions as well as in public agencies. They still required dedicated facilities and a team of operators dedicated to keep them running, but now they branched out from being numerical calculators into information

storage and manipulation devices. These computers could store and manipulate text as well as numbers. At the time computers were seen as the device itself and the giants of the industry were manufacturers whose logo was stamped on the hardware. The value of the unit was the hardware: by comparison the residual value placed on the software was almost incidental.



The evolution of the transistor into the integrated circuit at first enabled the mainframe to increase its capability. These systems became faster, and stored more information, but were still housed in dedicated facilities and operated by specialists. But the integrated circuit also allowed a useful processor to be squeezed into a single chip, and the computer industry was rocked by the advent of the personal computer. Few of the giants of the mainframe era made the switch, so the once prominent

brands of Univac, Digital and Burroughs quietly faded away. The advent of the computer as a consumer item changed this industry into a volume-based industry, and the change led to the emergence of new industry behemoths, and the most notable of these was Microsoft. Microsoft did not manufacture hardware. The company was a software house. While the personal computer hardware industry drove itself into price wars with razor thin margins due to unrelenting intense competitive pressures, Microsoft managed to establish a de facto monopoly over the software suite that turned these devices into essential components of almost every office on the planet.



The further evolution of the integrated circuit into ever smaller form factors with ever smaller power consumption levels and smaller heat dissipation requirements allowed the computer to shrink further to a physical package the same size as a human hand. The first mass market offering in this space, Apple's iPhone, was revolutionary in many respects. It reduced a general purpose computer to a device with just 5 physical buttons, not 101. It had a finely crafted colour screen of such a fine resolution that it stopped being a clunky computer screen and presented an interface that was natural to the eye. While the phone world and the computer industry had been interacting for many years, the phone world had jealously guarded its own territory, and while it

embraced digital technologies for transmission and switching functions, the telephone handset itself was slow to evolve from its basic speaking functions. It seemed that after more than a century the only innovation that the phone world came out with was the fax. The iPhone was a direct assault on this

staid, conservative world of telephony. One view is that the iPhone was a mobile phone that did so much more than send and receive calls. Perhaps a more realistic view was that the iPhone was a fully functional mobile hand-held computer that could incidentally be a mobile phone as well.

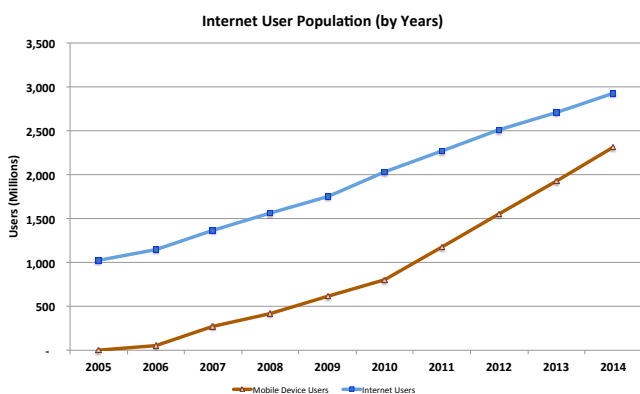
As has often been observed the path that computing has followed has been a path of continual disruption and innovation. Each wave of technology has created opportunities for adventurous entrepreneurs to unsettle the established incumbents, and the incumbents have proved to be incapable of keeping control of the further evolution of this same technology, creating windows of opportunity for further innovation and disruption. Change seems to be a constant factor in this environment.

Over the last 10 years the change that has been heralded by these mobile smart devices has been truly phenomenal. The computer is no longer in a dedicated building, and not even the centrepiece of a shrine on a dedicated work bench. What we see of the computer is a handy little device that we carry in a pocket, or even wear on our wrist. We expect the Internet to be always with us wherever we are and what ever we are doing. Paper maps? Who needs them when a navigation app is on my smartphone! Books? No, use your reading app! Messaging, talking, socialising, working. It's all on your smart phone. The social change is massive. The internet is no longer a destination. You don't need to go to a computer on a desk with a wired connection to some ephemeral network. It's with you all the time, wherever you are and whatever you are doing. The Internet is now incidental, but at the same time it's a constant background to everything we do.

The extent of this change can be seen in some industry statistics.

## Numbers, Numbers, Numbers

According to statistics published by the ITU-D the Internet now has some 2.9 billion users, or an average of 40 users per 100 across the world. As large as that number sounds, it is still dwarfed by the numbers of mobile phone service users, which today numbers 7 billion. These days in the developed world the penetration of SIM based services is now some 20% greater than the population. Many folk now have a mobile phone, a mobile tablet, possibly a SIM in their car, and potentially they have distinct mobile devices associated with their personal and professional lives.



In the developed world the number of mobile smart devices is now at 84% of the population, while in the developing world it is still somewhat of a luxury item with a penetration rate of 21%, but the total population of users of these mobile devices is now at 2.3 billion users of this mobile Internet. All of this growth has occurred within the last eight years, with an annual growth rate on these smart devices being activated at a rate of some 400 million per year, or 13 new device activations every second. Today these mobile devices now register as 40% of “visible” devices

on today's Internet, while the visible usage level of desktop and laptops has declined to 60%.

The supply side device production industry has also changed gears, and laptop and desktop processor manufacture is declining in the face of the uptake of mobile systems. In 2014 some 1.5 billion mobile smartphone units were shipped, and this quantity has had a massive impact on unit cost of such assembled devices and their component chips. The cost of an assembled smart device unit is now averaging well below \$100 per unit at the point of manufacture. Device costs are also contained through the use of open source software to control the platform, and of course the leverage of the

existing web-based universe of content to populate the device with goods and services means that the incremental cost of providing services to these devices is so small as to be incidental.

Who supplies this industry? In 2014 Google's Android operating system was used on 84% of all shipped smart phones and tablets, while Apple's iOS system was used for 12% of all devices, and the remainder was predominately Windows for Mobiles, which was mostly on the Nokia Lumina platforms.

These proportions of market share do not necessarily translate to equivalent ratios of revenue and corporate value. In looking at corporate value, Google's market capitalisation was some \$368 billion at the end of February 2014, which was comparable to that of Microsoft's \$359 billion market capitalisation. While Apple had just 12% of the share of shipped units in 2014, Apple's market capitalization is now \$755 billion.

The story as to why these market share numbers and capitalization levels are so different lies in the revenue figures from each of these companies.

The open software approach used by Google with its Android platform has not generated the same level of value for Google on a per device basis. Google are still strongly reliant on its advertising activity, which is reported to have generated some 90% of its total revenue of \$66 billion. It makes little in the way of margins from its hardware and software platforms in the mobile space, and the rationale for the open distribution of the Android platform may well lie in the observation that the resultant open access environment is readily accessible to Google's ad placement activities.

Microsoft's revenue of \$86 billion is dominated by commercial licensing activities, which appear to be a legacy of their historic position of dominance in the office IT environment. The Windows Phone revenues of \$2 billion had a margin of just \$54 million.

Apple is a completely different story. Apple's mobile products generated some \$150 billion in revenue, and. It should be noted that Apple has maintained a very tight level of control over both its hardware platforms, the iOS operating system, the software applications that run on this platform and the forms of interaction that these applications can have with the users. This form of tight vertical bundling, from the underlying hardware platform right through to the retail of services through the Apple Shop, has meant that Apple has been highly effective in maintaining very high margins on its iPhone products. Apple's legendary obsession with design in this space has meant that Apple has been able to maintain its position as a premium product in the eyes of the consumer. At the same time, the high levels of control over the platform has meant that Apple has been effective in realising this premium value as revenues directed to Apple. In so doing, Apple left little residual value accessible to others, including the mobile network operators.

## Mobile Carriage Technology

Data services in the mobile world have been nothing less than revolutionary over the past decade. The original model of data services in mobiles in the GSM architecture was to use the voice channel for data, and available data rates were typically 16 – 32Kbps, with an imposed latency of half a second.

The subsequent adoption of 3G standards headed in an entirely different direction, adopting the W-CDMA model with shared channels. Shorter transmission time intervals, an altered contention resolution algorithm, and use of amplitude and phase modulation techniques on the carrier signal were coupled with opportunistic channel acquisition. The result was a theoretical performance of 20Mbps downstream and 5.87Mbps up, but this theoretical speed is rarely achieved in practice. Speeds of up to 1 Mbps are a more typical user experience of 3G networks.



3G networks use a channel model between the gateway and the device, and use PPP to create the IP binding. One of the side-effects of this model is that each device association is bound to an IP instance. This implies that in order to support both IPv4 and IPv6 the operator would need to devote 2 channels to the dual stack device. The 4G LTE model strips out the PPP channel and the gateway and makes use of an all-IP internal infrastructure. The 4G device is now able to run in dual stack mode with both IPv4 and IPv6 associations to the network without doubling up on connection overheads. 4G also makes use of more sensitive digital signal processors in the handset and the gateway to support 64 QAM, and also uses MIMO multiple antennas to opportunistically obtain larger spectrum segments if they are available. On a good day, with the wind behind you, close to an otherwise idle base station with no interference it may be possible to drive a 4G connection to speeds of 100Mbps or higher. In practice most 4G connections run a lot slower than this, but it is still a visibly improved experience over 3G for most users.

While the mobile network operators are engaging in projects to select which particular variant of the family of (naturally) mutually incompatible 4G services they will use to upgrade their base station infrastructure, the technology is not standing still. At the start of 2015 Ericsson laid claim to the 5G label, claiming lab tests of a 15Ghz bearer, using the suite of QAM modulation and MIMO techniques to squeeze 5Gbps out of a radio connection. Ericsson claim that this would not be available as a general use technology until 2020, though I suspect that there are more issues besides this protracted lead time. The urban built environment is largely transparent to radio frequencies in spectral frequencies in the 100Mhz area, but as the frequency increases the signal has greater issues with absorption and reflection, and by the time you get to frequencies in this Ku band it is typical today to have a clear line of sight path. One suspects that much of the challenge for Ericsson and this 5G program will be to craft signal processors than can cope with the challenges of managing high frequency signals in a built environment.

## Futures

Where is the mobile industry headed?

The mobile phone market has already reached “super saturation” levels with the use of mobile SIM cards in the developed world exceeding the population. It certainly looks that the smart device market is heading in precisely the same direction. There are, however, a few cracks that are visible in this complacent view of the mobile industry’s future.

The underlying fuel for the mobile industry is radio spectrum space. There are two types of spectrum. The first is the traditional medium of exclusive use spectrum licenses, where the mobile network operator pays the government a license fee for exclusive access to a certain spectrum allocation in a particular geographic locale. The model of distribution of this spectrum has shifted from an administrative allocation model to one of open auctions, and the auction price of these licenses has, from time to time, reflected an irrational rush of blood to the heads of the network operators. The high cost of spectrum access implies that the network operator starts with a non-trivial cost element for the spectrum, and on top of this the operator must also invest in physical plant and business management operations. Typically, the spectrum actions are constructed in a manner that there is no ability for a single operator to obtain a monopoly position, and most regimes ensure that there are between two and four distinct spectrum holders in the most highly populated locales. Sometimes this level competition is enough to maintain an efficient market without price distortions, while at other times the small number of competitors and the barriers to entry by any new competitors leads to various forms of price setting distortions in the mobile market, where the retail price of the service has no direct relationship to underlying costs.

But competitive pressure from other mobile network operators is not the only source of competition. The other form of spectrum use is also a factor. WiFi systems use two unmanaged shared spectrum

bands, one at 2.4.Ghz and the second at 5Ghz. There are typically limits on the maximum transmission power used by devices that operate in these bands, but to all other extents the spectrum is effectively open for access. For many years WiFi has been used to support domestic and corporate access. WiFi systems typically operate within a range of up to 70m indoors and 250m outdoors. This small radius of WiFi systems, which is an outcome of the limited transmission power, has fortuitously also allowed these system to operate with extremely high capacity. They run at speeds that range from 10 to 50 Mbps (the 802.11b specification) through to speeds of up to 1.3Gbps (the 802.11ac specification), and expectation is that this can be lifted to even higher speeds in the near future.

For a while these two spectrum use models have compartmentalised themselves into distinct markets. Exclusive-use spectrum for the 'traditional' mobile network operators and shared spectrum use for self-installed domestic and office applications. But now some operators of wired network access infrastructure are heading into head to head competition with the mobile operators in the provision of mobile services. Comcast's Xfinity service in the US is a good example of this approach, and it boasts of millions of WiFi hot spots which are usable by existing Comcast customers without any additional cost. Perhaps surprisingly, some of the incumbent mobile operators have followed this lead and are also offering free WiFi access services to their customer base. For example AT&T has a WiFi offering in the US. Part of the rationale for this may well be protecting their market share. However there is also the ever-present issue of congestion in the licensed radio spectrum space. 3G and 4G data services are opportunistic, and scavenge otherwise uncommitted access capacity. However, in places of intense use, such as high density urban centres, the challenge of providing high capacity data services in licensed spectrum becomes a tough challenge. One approach is to use WiFi access points as a relief mechanism for these areas of otherwise high congestion. The observation that this clashes with the business objectives of the mobile operator's WiFi-based competition is perhaps a convenient bonus in such a scenario!

But there are some issues with WiFi that are not as obvious in the traditional mobile service space. The mobile industry has supported base-station handover since its inception, so that an active data stream to a device can be supported even as this device is in motion, being handed off from one radio access point to another. WiFi handoff is not as cleanly supported. The issues surface when the WiFi access points reside in different IP networks, so that a handoff from one WiFi access point to another implies a change of the device's IP address. Conventionally a change of IP address equates to a disruptive change to all of the devices active connections. However, it is possible to create an application that uses some form of session key persistence to allow one of the end points of the session to change IP addresses while keeping the session open and not losing state. We have also seen mobile applications make use of Multi-Path TCP, where a logical TCP session is shared across multiple interfaces, potentially permitting interfaces to be added and removed from interface set supporting the logical TCP session. In essence, the device and its applications are now able to leverage a rich connectivity environment where all the interfaces in a mobile device can be used as appropriate.

What does this mean for the mobile industry?

What we are seeing is that the mobile device is no longer tethered to a mobile network operator, and the device is able to react opportunistically to use the "best" network, whether it's the greatest available capacity or the lowest incremental cost to the consumer. From the device's perspective the mobile network is just one possible supplier of transmission services, and other options, including WiFi, Bluetooth and USB ports can also be used, and the device is able to make independent choices based on its own preferences.

This has profound implications. While the device was locked into the mobile network, the mobile network could position itself as an expensive premium service, with attendant high prices and high revenue margins. The only form of competition in this model was that provided by similarly positioned mobile service operators. The limited number of spectrum licences often mean that the players established informal cartels and prices remained high. Once the device itself is able to access other access services, then the mobile data network operators find it hard to maintain a price premium for

their service. The result is that mobile service sector is being inexorably pushed into a raw commodity service model. The premium product of mobile voice is now just another undistinguished digital data stream, and the margins for mobile network operators are under constant erosive pressure. The unlicensed spectrum open WiFi operators are able to exert significant levels of commercial pressure on the mobile incumbents in the mobile service environment. This means that the premium prices paid for exclusive use spectrum licenses are exerting margin pressures on operators whose revenues are increasingly coming from commodity utility data services.

Perhaps there are yet more changes on the way. The “improvements” that turned the mobile phone into a smart mobile weren’t motivated by providing a better voice experience. Not at all. What changed was that the device added retinal displays, cameras, touch inputs, local storage, positioning, and access to all of the online services that we were used to with desktop computers. What we have today is the equivalent of yesterday’s general purpose mainframe computer in your pocket. The mode of improvement in this model is more memory, sharper displays, better power usage, improved processing capability. In other words marginal improvements to the same basic model. But we are also seeing a new range of specialised products that also use mobility but are not general purpose computers, but dedicated devices intended to provide a particular service. The “Fitbit” is a good example of this form of specialization, which is a specialised device that measures the wearer’s physical activity. Credit cards, and many municipal metro cards, now include embedded processors that store value on the card itself. Many countries now issue passports with embedded electronics.

So will the model of a general purpose computer persist into the future, or will we see further specialization as the computer industry abandons general use models and instead specialises and embeds itself into every aspect of our lives?

It seems that we are indeed seeing the Internet, and even the computers that populate the Internet slowly fade from public prominence. It seems that these technologies are so much a part of our lives that we no longer recognise them as something special or distinguished. They are inexorably weaving themselves into the way we live our lives. And the crucial element of this transformational change has been the untethering of the Internet and its enthusiastic adoption of mobility. Once the Internet is always present and always available it no longer is visible. Like the water from your tap, we will only really be aware of the extent of the role of this ubiquitous internet will be in those rare times when its just not there!

---

## Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

---

## Author

*Geoff Huston* B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

*[www.potaroo.net](http://www.potaroo.net)*