NB-IoT is Dead. Long Live NB-IoT.

We're at the point where IoT battery life is greater than the comms network life.

As the old adage goes, "while the cat's away, the mice will play". In the case of NB-IOT, "when the spec's delayed, LPWAN will play", which is exactly what's happening in the Internet of Things market today. The problem is that 3GPP (the 3rd Generation Partnership Project), the standards body which has been responsible for the 3G, 4G and 5G mobile standards, dropped the ball as far as the Internet of Things is concerned. Seduced by the slabs of black glass which suck up both our attention and the mobile networks' spectrum, the 3GPP engineers totally forgot to design something to replace the old 2G workhorse of GPRS, which is responsible for most of today's machine to machine communications. Instead, they spent all of their time designing high power, high speed, expensive variants of 4G to support an ongoing dynasty of iPhones, Galaxys and Pixels, none of which were any use for the Internet of Things.

Noticing this hole, a number of companies who had been developing proprietary, low cost, low speed, low power communication options saw an opportunity and created the Low Power WAN market. Whilst many perceived them as a group of Emperors with no clothes, the network operators were so desperate to have something to offer for upcoming IoT applications that they started engaging with them, rolling out LPWAN infrastructure. Whether they believed the LPWAN story, or just hoped it would fill a hole is difficult to ascertain, but no-one can deny that LPWAN is now firmly on the map, in the form of Sigfox, LoRa, Ingenu and a raft of others. To address that challenge to their hegemony, the GSM Association (GSMA) directed the 3GPP to assemble their own suit of imperial clothing which would be called the Narrow Band Internet of Things, or NB-IoT.

This is the story of why NB-IOT was too late, why it will fail in the short term, why it will win in the long term, and why the industry will struggle to make any money from it.

One of the most surprising aspects of this story is how long it took 3GPP and the network operators to realise that they had a problem. It's not as if they didn't see the problem coming. Back in 2010, Ericsson set the bar for much of the subsequent hype around the Internet of Things by making a very public prediction that by 2020 there would be 50 billion internet connected devices. They've subsequently downgraded that, but very few in the industry noticed – for them, it's very difficult to discard the prospect of "tens of billions" once it's made its way into their business plans. Numbers that big get attention in boardrooms, whether or not they mean anything – they just sound so good that they are assumed to be true.

What happened is that the industry became fixated with the concept of revenue today, rather than revenue tomorrow. As users embraced smartphones, their demand for data soared. When competing smartphone vendors made smartphone screens larger, mobile video took off, putting further pressure on the network's capacity. Everyone's attention became focused on how to build enough capacity into their network to retain their users. Instead of calling for new standards for M2M and IoT, operators started concentrating on how they could use their existing spectrum more efficiently. There was an easy answer to this – turn off their old 2G networks and use them for 4G, which supported around 40 times as many users. It was only as they started to do this that they belatedly realised that they were euthanising the only technology they had which would support the Internet of Things. At which point the LPWAN industry stepped into the frame and started cutting deals. The GSMA panicked, and directed 3GPP to embark on the path to NB-IoT.

At this stage it's worth pointing two things out. The first is the normal timeline for developing a new radio standard, and the second is the requirements for the majority of the projected 50 billion IoT devices.

Developing a wireless standard is a slow business. Back in 2010 I tried to estimate the time and cost involved and came to the conclusion that it costs around a billion dollars and takes 8 – 10 years before the standard is robust and getting traction in the market. That was for personal area networks like Bluetooth, Wi-Fi and ZigBee. Cellular networks are more complex, so cost more and take longer. Despite the evidence, the GSMA announced that their new NB-IoT standard would be complete and released in six months. Six months later, they announced that it was going well and that they would release it in six months' time. And six months after that they put out a press release saying that the specification was complete. We'll come back to that in a minute.

The second thing we need to look at is what a standard for wireless IoT connectivity needs to do? Most IoT devices will be quite taciturn. They will measure data and events and send that data a few times each day. They're not going to be streaming video or having lengthy conversations because they're battery powered. If they're going to run for several years on a small battery or some energy harvesting power supply, all they can manage is a few messages each day. Sigfox understand this and make it evident in their data plans. They're not talking about hundreds of Megabytes like the cellular industry, but as little as 14 messages of 12 bytes each day. That's about the same as a single SMS message. To put it another way, most IoT applications make text messaging look bloated.

It's not at all clear that the GSMA understand this. In a recent Mobile Broadband Forum meeting, the GSMA and other operators kept on implying that IoT devices need data rates of tens or hundreds of kilobits per second. That is definitely what network operators want to sell, but it's not what IoT devices need. If we're going to get to billions of device, connectivity and silicon needs to be cheap. Cheaper and simpler than GPRS was. The cellular industry has never taken on board that fact that the reason we don't already have billions of IoT devices is that even GPRS is too expensive. Trying to make NB-IoT more complex than GPRS is not going to kickstart the IoT era. What we need is a standard which will let companies make a chip that costs around a dollar in high volume.

That's not where the cellular chip industry has been going. In the early days of 2G, networks operated at two different frequencies, with relatively simple radio modulation. That meant that chips were moderately simple. Over time, the GPRS modules which are used in most current IoT devices have fallen in price to around $7. However, as the desire for more bandwidth has grown, 3G and 4G chips have become much more complex. Moore's law has helped to prevent them becoming exorbitant, but each new release of the standards has to support a growing number of frequency bands (we're up from 2 to over 70), as well as all of the different protocols in the previous standards

which have gone before it. Developing these is prohibitively expensive. As a result, 3G modules cost around $20 and 4G modules $35. The growing complexity, which requires immensely complex protocol stacks to complement the chips, has benefitted a very few silicon suppliers, who have largely destroyed the competition. Qualcomm dominates, with Mediatek taking most of the rest of the market. The business model for both is to sell billions of chips to a small number of high volume manufacturers who have deep technical competence to integrate these into their products. That is very different to the model needed to support tens of thousands of IoT manufacturers who need $1 comms chips which they can just drop into their products.

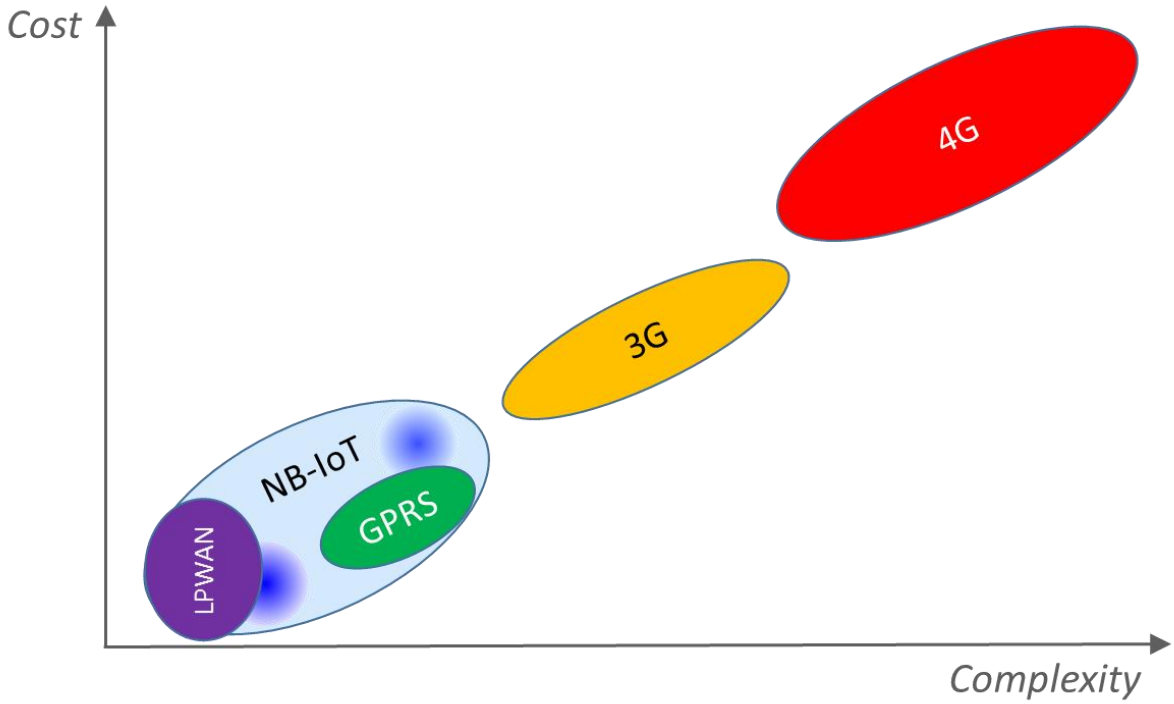

You can see this contradiction in the NB-IoT standard which has recently been released. There were two industry groupings with radically different approaches. The traditional one, led by Nokia and Ericsson, proposed what is essentially a cut down, lower power variant of 4G. The key feature of this is that it is capable of working with other 4G devices in the same spectrum, so it can easily be slotted into existing networks. However, to do that it needed to retain a fair degree of radio complexity to be aware of other 4G traffic. That has two consequences. It meant the chip was much more complex because it had to be able to identify what was going on around it, hence it's still expensive. It also made it more difficult to make it very low power.

The alternative approach, led by Huawei and Vodafone was for a "clean sheet" approach. This was a solution which did not have the intelligence to coexist with 4G networks, but required operators to set aside a small amount of spectrum for it, (which could be a guard band), specifically reserved for IoT traffic. As the chips didn't need to be aware of any other 4G traffic, they could be much simpler and hence much cheaper. It's a cleaner approach, but one which goes against the traditional network approach of making complex hardware which can work on any band around the world. Network operators typically prefer the complex hardware approach, as it passes the problem of global interoperability onto the chip and protocol stack companies. Whatever the operators do with

their networks, regardless of the frequency bands they own, things just work. But it raises the cost of hardware.

This "clean sheet" approach grew out of the Weightless standard. Neul – a Cambridge start-up helped developed Weightless as a new radio and protocol for use in TV Whitespace. That failed to get traction, but the company was acquired by Huawei and the technology repurposed to work in the licensed spectrum that's used for LTE. Because it does not have the baggage of backward compatibility, there's a fair chance that the silicon could get down to the $1 mark.

These two approaches are essentially incompatible, and it was interesting to speculate how 3GPP would resolve the difference between them. Hence I was intrigued to see the resulting specification when it was published. When you start to read it, you can see how they managed to get it out so quickly. Instead of trying to find a compromise, it includes both the Huawei / Vodafone and Ericsson / Nokia / Intel options, so it is entirely up the chip vendor and network operator to decide which they support. That means that a user or manufacturer has absolutely no idea of whether an NB-IoT product they make will work on any particular NB-IoT network. It's as if the acronym should really be Nobody Believes the Internet of Things.

It's a fudge, where the specification group has produced some pieces of paper to meet a deadline and then passed everything over to a PR department which is taking the post-truth approach to promoting the technology. It would be nice to think that the specification group had realised that this first release was just a PR exercise and were working on harmonising the two conflicting proposals, but it seems they're ignoring that and looking at adding location features instead, presumably because LoRa is offering that, and they don't want to be left behind again. In other words, bells and whistles are more important to them than making NB-IoT work.

Making it work appears to be left to market forces. Vodafone is trumpeting the first commercial NB-IoT network. At the same time, Sonera, in Finland is announcing the first commerical NB-IoT trial. Although that may seem confusing, there is no contradiction here. Both are telling the truth, as Vodafone is using Huawei's NB-IoT, which is totally different for the Nokia NB-IoT which Sonera is using. Nobody knows which variant will win. The key player in this could end up being Huawei. They have a captive silicon supplier in Hisilicon, which should help them get to the $1 chip price point. If they could persuade the Chinese Government to deploy hundreds of millions of devices in the country, this could make it the de facto standard. Nokia, Ericsson and Intel are unlikely to concede without a struggle, but with a higher cost and the lack of scale that a Government backed deployment in China could provide, they may struggle to gain momentum.

Unfortunately, this type of commercial battle generally doesn't help the market. Without global compatibility, manufacturers will be loath to adopt the technology, as they have no idea whether it will work in any target market. That reduces volumes, which keeps chip costs high. It also delays all of the important things like developing test equipment and compliance programs which are vital to develop a robust network, which further undermines confidence. To survive, NB-IoT needs to be a single low cost, globally interoperable standard. In its current form, NB-IoT is dead.

While it goes through its death throes, the LPWAN suppliers will make mischief.

Sigfox is being aggressive in pricing, both for modules and data contracts. They recently announced that modules will be available for just $3 in 2017 and already have data plans with charges as low as $1.50 per year. They also desperately need to get the number of connections up, so will probably offer even lower costs in the near future. The company has raised over $300 million in funding and is aiming for an IPO in 2018. However, they feel that they need to get above 100 million active

devices to persuade the market to support a decent valuation. So their investors will be putting pressure on them to get more connections made as soon as possible, potentially commoditising the IoT connectivity market in an attempt to buy market share from their rivals.

LoRa is a more distributed community, with multiple vendors providing parts of the ecosystem. However, LoRa has a significant difference from other LPWAN offerings, which could be important. It is the fact that anyone can buy a gateway and set up their own network. A crowdfunded initiative – the Things Network, has designed modules and gateways and persuaded the electronics distributor Farnell / Element14 to sell them in the same way they sell Raspberry Pis. For those who don't know it, the Raspberry Pi is a highly effective embedded computing board. Originally designed to help teach coding in schools, it has been adopted by the maker community as the basis for thousands of projects and products. Farnell have recently announced that they have shipped their ten millionth Raspberry Pi.

The Things Network / Farnell initiative is relevant, as they will be selling LoRa gateways for €250. In other words, for €250, anyone can become an Internet of Things network operator covering a radius of around 5km. The Things Network - a development community attempting to build a global LoRa network, is providing compatibility layers behind that which will stitch many of these gateways together. Costs will probably be slightly higher than Sigfox, but this will appeal to an open source community, with the innovation benefits that brings to an emerging technology.

There are issues about scaling. Tech hotspots like Cambridge, Amsterdam and Berlin could each have over a thousand LoRa gateways by Christmas 2017, which could make or break the technology. It will be an interesting experiment. It may also give Ingenu an opportunity, as they've been in the game longer and appear to have a more robust technology in terms of scalability. But they've not achieved the same traction in the minds of IoT developers yet.

This brings us to the important part, which is what this means for network operators? Other than Vodafone, who have firmly nailed their colours onto the NB-IoT mast, most operators are hedging their bets by flirting with at least one proprietary LPWAN option. However, in order to get critical mass, contract prices are racing to the bottom. SK telecom is down to $0.30 per month and Sigfox's pricing will probably push that down to below $2 a year in the near future. That's a long way away from the $50- $200 that operators get from their current M2M contracts.

At $2 a year, 20 billion devices will contribute around 4% of current global mobile subscription revenues. That is probably less than network operators currently make from their GPRS subscriptions, yet it will replace much of that revenue. In other words, by supporting 20 billion IoT devices, the network operators will probably be making less money. Let me emphasise that point. The IoT opportunity of tens of billions of connected devices could reduce mobile operator revenue, not increase it.

Many mobile operators seem to think that they will make money from other parts of the IoT value chain, like cloud services or data analytics, but there is little indication that they're well positioned for that. Amazon, Google and a host of others are already there. In the next few years, the volume in deployments will probably be using the LPWAN standards of Sigfox and LoRa. The developers who choose them will naturally turn to Amazon and Google, giving them the opportunity to further refine their IoT offerings. I'll cover this in more detail in a future article.

Despite the present debacle over NB-IoT, the developers at 3GPP are bright – they will eventually get a specification out which meets the industry's requirements, whether that's driven by market forces winning out or technical decisions. However, my guess is that it may not be before 2023, as

that's how long wireless standards take. Which gives the different LPWAN standards plenty of time to play, and time for the cloud and analytics providers to shake out, settle down and start some serious customer acquisition.

The great thing about 3GPP standards is that they're dead easy to roll out. In most cases they're simply a software upgrade for the base stations. So it won't take long to go from a final standard to global availability. At which point most IoT manufacturers will probably migrate to it, signalling the end of the short-lived LPWAN era. Of course, most of the LPWAN players and their investors are looking for shorter term returns, so they may already have disappeared. Even five years is a long time in a venture funded world.

What will be missing in the future NB-IoT world will be the hoped-for revenue. The years of LPWAN competition will have driven any profits out of NB-IoT, leaving the operators as pipes. It will also have established other players higher up in the value chain who can cream off what profit there is to be made. A future variant of NB-IoT will come to life and dominate as the connectivity standard for IoT, not least because as volumes grow, the licensed spectrum that operators own will offer a Quality of Service that is missing from the LPWAN offerings. It will also provide the certainty that manufacturers are desperate for, which is that the network will be a stable solution which is available for fifteen to twenty years. NB-IoT will wipe out any remaining alternatives, but it will not be the IoT pot of gold that many in the industry believe.

There is a final sting in the tail of this story, which is that for years we have been striving to develop low power, wide area connectivity which will enable a sensor battery life of ten years or more. The irony is that we now have a set of different LPWAN options which look as if they do support a ten year battery life, but it's unlikely that any of them will still be operating in ten years' time. In other words, battery life now exceeds network life.

One wonders how we got to this point? There is little good news for an equipment manufacturer, who is faced with the prospect that whatever connectivity solution they choose today, it will probably disappear within the next ten years. In other words, their product obsolescence is in the hands of their choice of network operator. But that's the problem when you forget your King is dying and everyone spends their time running around backing pretenders to the throne. Be careful what you wish for. NB-IoT is dead. Long live NB-IoT.

Nick Hunn

December 2016

nick@nickhunn.com

+44 7768 890 148

You can read my previous articles about NB-IoT, LoRa, Sigfox and the IoT on my Creative Connectivity blog.