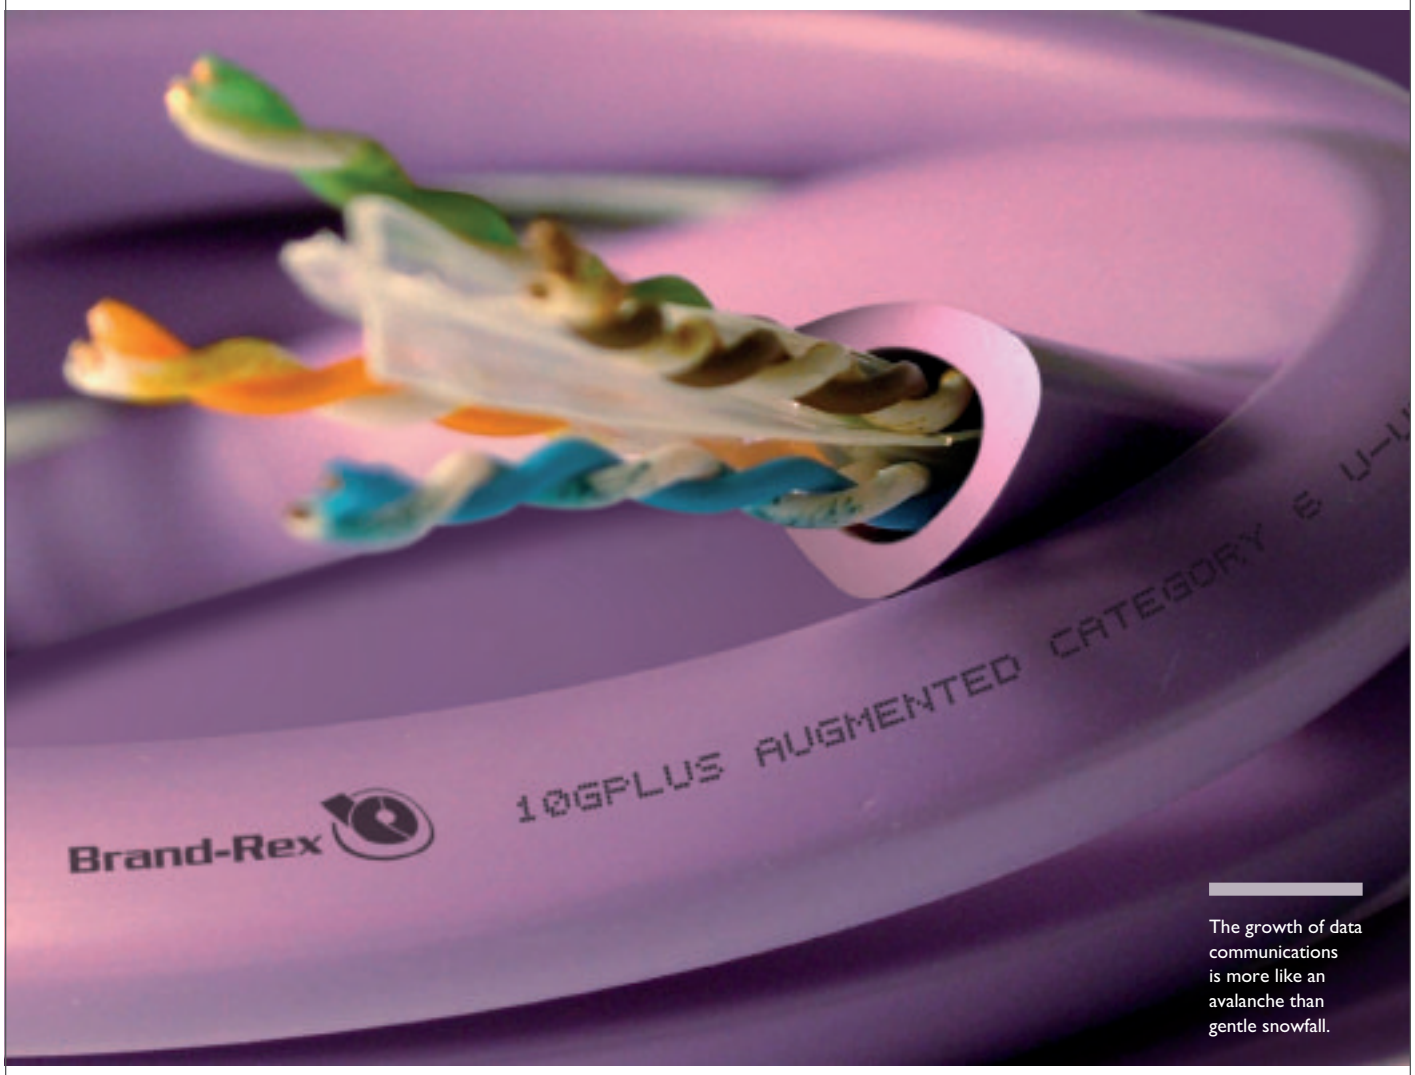

The growth of data communications is more like an avalanche than gentle snowfall.

Into the unknown

Do we need 40Gb and 100Gb Ethernet over copper? In the data centre; the answer is an unequivocal yes. **Ken Hodge** of **Brand-Rex** looks at some of the drivers, the state of current knowledge regarding these possible new extra high speed copper Ethernet technologies and gives some sage advice on what users and installers should be installing in 2012.

The world we live and work in is, as we know, becoming more and more 'online'. In our private lives it's Facebook or Twitter, online gaming and video on demand streaming services like Netflix or Amazon's Lovefilm sending multiple GigaBytes per film over our home broadband to our PC, X-Box or even to our smartphones.

Data explosion

In business we're increasingly using hosted services like Google Docs or SalesManager.com, webcasts, LinkedIn, Facebook, Twitter and enterprise scale hosted apps such as enterprise resource planning. The mobile phone has morphed from a simple analogue voice device to a mobile computer, video, calling device, online TV receiver that almost every man, woman and child carries with them 24/7. The volume of data that these devices threaten to produce and consume is astronomic.

If you need any further convincing that the growth of data communications is more like an avalanche than gentle snowfall, there is the inevitability of massive amounts of machine-to-machine (M2M) data communication in coming years.

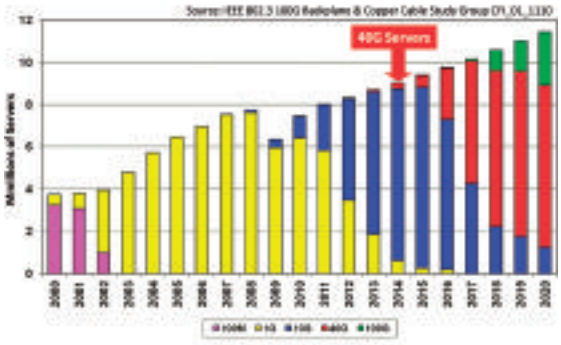

Servers by Ethernet port speed

In the data centre specifically, the move to virtualisation – with dozens of virtual servers running on each piece of server hardware means that each physical box's Ethernet I/0 is now handling the aggregate traffic of all of those virtual servers which in the old days would each have had their own 1Gb/s Ethernet links. In the backbone too, users are already aggregating multiple 10Gb/s links to achieve higher effective speeds because 10Gb/s already isn't enough.

And so in fact, it's not a question of 'do we need 40G and 100G?' – it's also a question of 'where do we go after those?'

Is 40Gb/100Gb over copper actually possible?

As far back as August 2009, Penn State University in the USA undertook modelling studies which showed fairly convincingly that up to 50Gb/s over 100m should be possible (meaning that 40Gb/s certainly could be). At 100Gb/s the study predicted a maximum of 70m.

Our own predictive studies at Brand-Rex came to similar conclusions with the caveat that the greater the distance required the greater the complexity and cost. So it's important to understand that the final solution will be based not just on physics but on economics with a fair injection of industry politics.

Do we need 100m links?

Historically, all of the 'categories' of cabling for Ethernet have been based on a 100m channel. That is because they were originally designed for the enterprise horizontal network connecting out to desks and workstations. But it's debatable whether there is sense in sticking to this dogma for 40Gb/s, 100Gb/s and faster since, realistically, they are unlikely to be used in the enterprise horizontal in the next 10–15 years.

Why does this matter? Well the longer the link length – the more difficult to design, complex and expensive is the cabling system and its electronics.

In the data centre today with our own Category 6A 10Gb/s zone cable - which we miniaturised by reducing its reach to 70m - we know that this covers approximately 97 per cent of all possible data centre links. And in fact, if you know you have a 70m limit when you design the data centre, you can easily design it so that all links are less than 70m.

Link Length	% of ALL data centre links covered
0-10m	33%
0-30m	80%
0-50m	93%
0-70m	97%

So, given that it is far easier to design the electronics and cabling system for a system capable of 70m or 50m – and that it is possible to design a data centre with all links 70m or less (and use fibre for any really long ones) this poses the question of whether 100m is a 'sensible' requirement this time around. 'But we've always done it this way!' is a very bad premise from which to start any new project!

The pattern re-emerges

A pattern has previously emerged which looks set to repeat itself with 40/100Gb/s. It's this: Fibre gets there first with the new higher speed. It's very expensive at first – but the cost is justifiable for those leading edge applications that need the new higher speed long before it becomes mainstream. And for 40 and 100Gb/s, we've already seen the fibre standards published and systems becoming available.

Fairly quickly on the heels of faster fibres comes short reach copper for equipment backplanes and short interconnects. That's because the electronics needed for short reach is vastly simpler than for long channels and more expensive high performance cable and connectors can be justified.

With short reach Gigabit Ethernet the limit was 25m on twinax cable. By 10Gb/s this was down to 15m and the recently defined standards for 40Gb/s and 100Gb/s it is 7m max. But it's useful for early specialist requirements.

And historically, a couple of years later along comes the new generation of Base-T. The cost of the copper channel soon becomes a mere fraction of that for both fibre and short reach copper. At which point the short reach copper becomes obsolescent and the fibre reverts to being a solution for longer distances.

I predict we will see exactly this same process with 40Gb/s and 100Gb/s – although they may both go through the process in parallel.

UTP at 40Gb

You might think, as the 10Gb/s market has now become skewed towards shielded and the issues of alien crosstalk are likely to be orders of magnitude worse at 40/100Gb/s than at 10Gb/s, that UTP would be out of the question. And many in the industry are saying so.

But I'm definitely not counting it – or maybe a hybrid solution – out of the running! I've seen recent papers from highly reputable engineers indicating that UTP could be possible even at these esoteric speeds. And for those who remember the early development of 10Gb/s; it was said by many to be impossible then too.

How will 40Gb/ 100Gb work?

With the exception of the 40Gb/s and 100Gb/s short reach copper solutions it is essential to understand that – despite the marketing claims of some manufacturers – no one can possibly know!

We in the R&D world all know lots of ways they could work. But these use many different combinations of cable characteristics (drawn from cable types that already exist and those which are not yet even designed) together with multiple options for line coding and digital signal processing (DSP). Dozens, if not hundreds of possible permutations. No one can yet predict which will win out technically, commercially and politically.

In fact – the IEEE working group in which these things will be thrashed out – usually over about a two year period – hasn't even been formed yet. It's looking likely that it might be sometime this year but even if it is formed it will be a long time before there is any hint of certainty.

Data centre server-to-switch link lengths

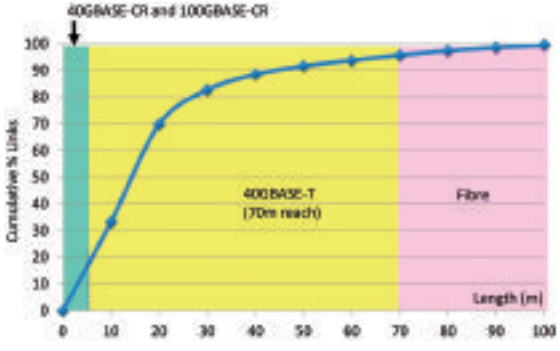

So – if any vendor is currently trying to sell you a '40Gb/s or 100Gb/s ready' copper cabling system be very, very cautious.

Can my existing cabling support 40Gb?

Following on from the above, the only honest answer is no one knows. For Category 6A it's pretty unlikely, likewise for Category 7. Category 7A is a slim possibility – but it's more likely to be a revised Category 7A (more of a Category 7AA perhaps) at 1,200MHz instead of the current 1,000MHz or perhaps a new cable at an even higher 1,600MHz. So again it's maybe or maybe not.

Our own modelling together with DeMontfort University using the Technology Forecasting Technique indicates that a 2,000MHz cable frequency could actually prove to be needed. It's going to be quite a bumpy ride while all of these options are fought out once the IEEE study group is formed and starts its work.

Will it be RJ-45?

So far the industry has managed to re-invent the RJ-45 connector at every stage of Ethernet's evolution. From 10Mb/s to 100Mb/s to 1Gb/s and currently 10Gb/s the RJ-45 has remained the connector of choice.

Other connectors have been designed along the way with potentially higher bandwidth capabilities – but with one exception (the ISO/IEC 60603-7-71) which combines conventional RJ-45 for up to Gigabit with 'four corner' pairs and a slightly complex in-connector switch for higher speeds, they are not backward compatible.

These two alternative connectors have never been deployed by the switch and server manufacturers who have stuck rigidly to RJ-45 which has always been backward compatible. Another of these connectors is the ISO/IEC61076-3-104 square connector but as this is thicker than the latest Apple laptops that pretty much rules it out from deployment on mass market user equipment and so it's unlikely that the main switch and server vendors would adopt it either.

Another connector – the ISO/IEC 61075-3-110 – has the same form factor as the RJ-45 our industry loves but has one pair in each corner for greater isolation. Unlike the 60603-7-71 it isn't backward compatible, but it may prove to represent a better option than the other two contenders from a commercial and political perspective.

But let me throw a ball in from left-field, as they say! I'm far from convinced that the dear old RJ-45 can't once more re-invent itself. And if the industry can squeeze the necessary performance out of it – even if it means a tighter cable specification – I reckon that such an RJ-45 could win hands down. But on that score – only time will tell.

What to install today?

'OK Ken,' you're probably thinking, 'now you've told me it's a minefield out there – what on earth do I install now?' My answer to that is very simple - Category 6A. It's going to do most things you'll need now and for the next 10 years. If you need faster, use short reach copper or fibre.

Category 6A is going to do most things you'll need now and for the next 10 years.

Yes, you could deploy Category 7A cable with Category 6A connectors as some vendors are recommending – but to be honest the likelihood of that being the cable that is ultimately chosen to support 40Gb/s or 100Gb/s in the standards are, at best, slim. That said, this combination will give you a decent Category 6A system, but it's not a short odds bet for higher speeds. And of course, if you're just looking for a short term installation such as temporary rented building space then Category 5e is still fine and will stay so for many years to come.

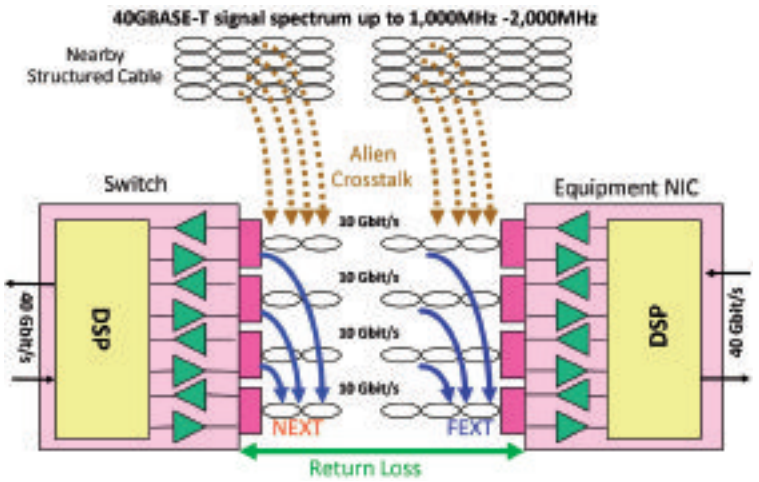

40Gbase-T will need four 10 Gb/s bothway lanes

Ken Hodge was educated at Liverpool University and Warwick University. He is a chartered engineer and a Fellow of the Institute of Engineering Technology. He is currently chief technical officer at Brand-Rex and is based in its headquarters in Glenrothes, Scotland.

Ken has been involved in the industry since 1982, researching, designing and developing optical fibre and high frequency cabling for LAN and telecom networking. He is actively involved in BSI, IEC and CENELEC activities in international standardisation in the cabling sector.