Implications of the next-generation Internet protocol for ABB

IPv6 the new Internet protocol

Niels D. Aakvaag, Nils A. Nordbotten

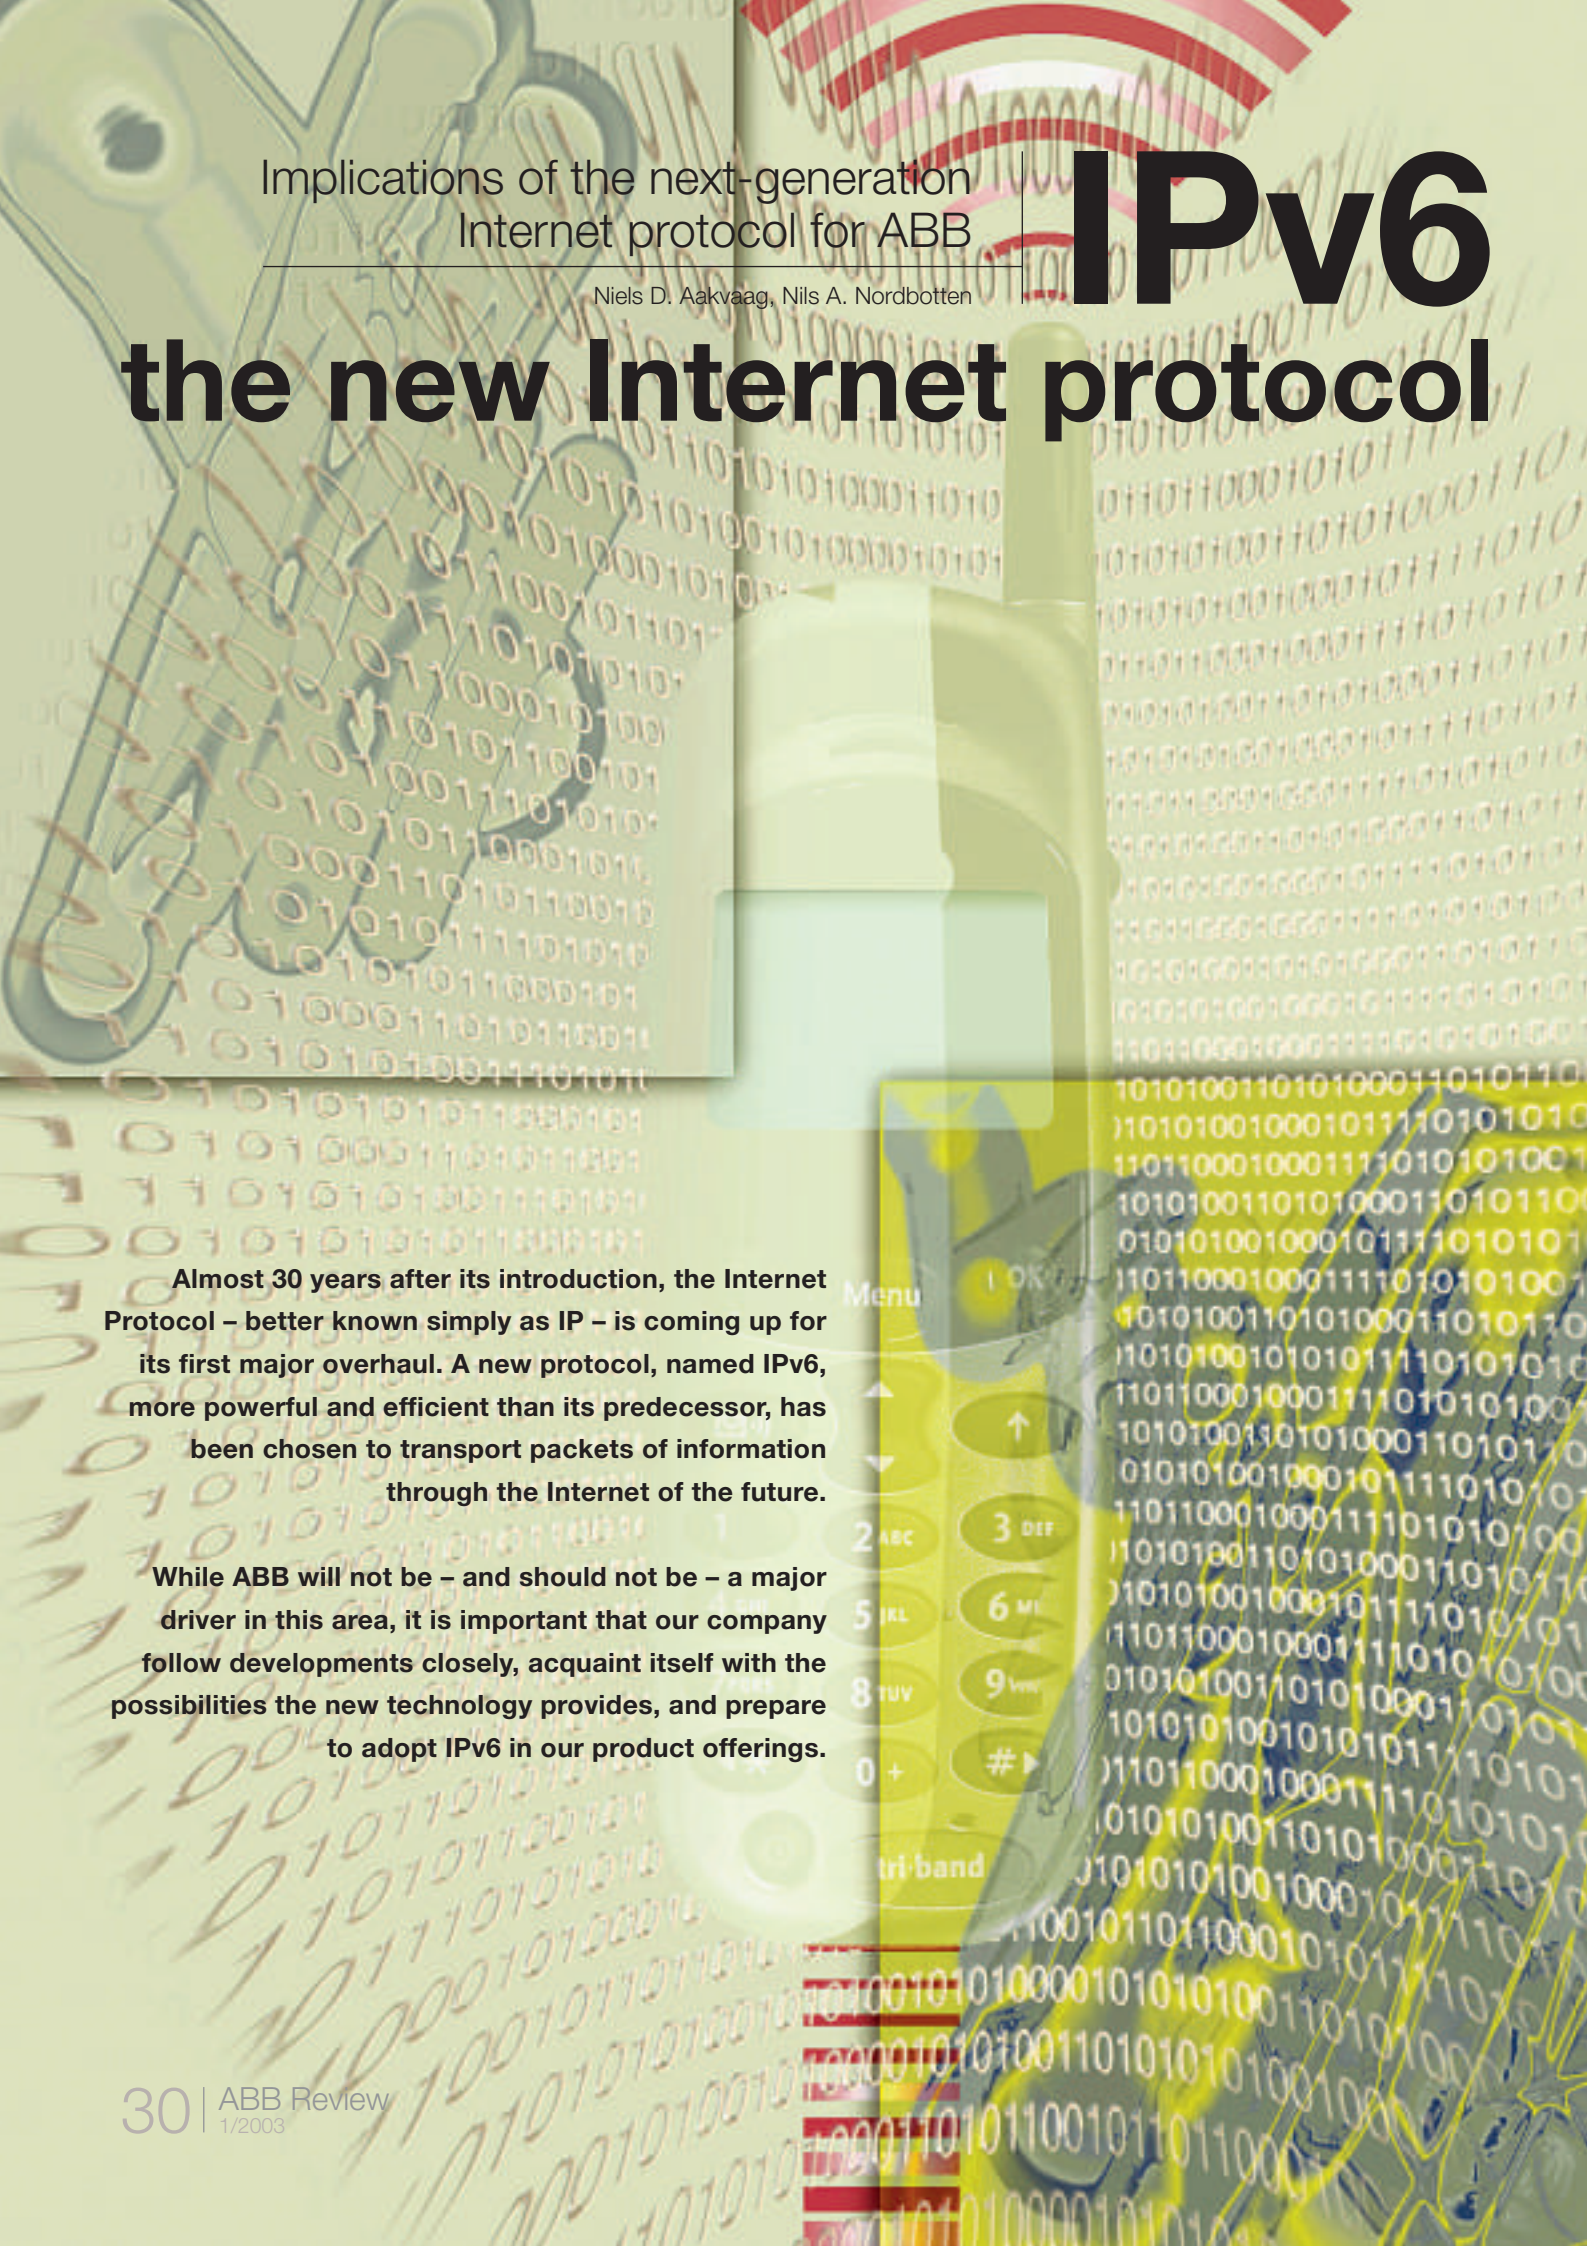

Almost 30 years after its introduction, the Internet Protocol – better known simply as IP – is coming up for its first major overhaul. A new protocol, named IPv6, more powerful and efficient than its predecessor, has been chosen to transport packets of information through the Internet of the future.

While ABB will not be – and should not be – a major driver in this area, it is important that our company follow developments closely, acquaint itself with the possibilities the new technology provides, and prepare to adopt IPv6 in our product offerings.

The Internet has been with us for almost 30 years, but it is only during the past decade or so that we have come to look upon it as an indispensable tool in our everyday lives. Today it is *the* information carrier – the provider of basic services on which we all depend, like email, voice data, and webcasts, besides providing the essential infrastructure for industrial communication.

The standard currently being used to transfer data on the Internet is known as Internet Protocol Version 4, or IPv4. Now nearly 20 years old, IPv4 has shown remarkable resilience. Its flexible structure has allowed new protocols to operate in conjunction with it to provide services like support for time-critical tasks, the ability to address people on the move, and secure communication – all tasks for which the Internet was never really intended. Time does take its toll, however, and changing user patterns, as well as the increase in traffic, are imposing new requirements on the Internet.

So what's new about IPv6?

Revised IP header

A comparison of the respective headers of the IPv6 [1] and IPv4 is a good starting point for any discussion of the former version's virtues 1. The most obvious differences are the IPv6's augmented address length and the way in which the remaining fields has been altered. Fields defined in IPv4 that have turned out to be of little use have been removed. For example, the 'Length' field has been removed (the IPv6 header always has the same length), all fields associated in any way with fragmentation have been removed (with IPv6, fragmentation is not allowed in the network, only at the end nodes), the 'Checksum' has been removed, and the 'Optional' field has been removed. The latter was included in IPv4 to cater for information on all imaginable functionality in the same header. In IPv6, this has been resolved by the concept of extension headers, where a new header is appended to the current header only when required. This is expected to yield a much higher processing speed.

Address space

What has driven the development of the IPv6 more than anything else is the growing shortage of Internet addresses – with the 32-bit address field in IPv4 only approximately four billion addresses are available. Aggravating the problem is the inefficient use of the available address space and the fact that it is unevenly distributed. The USA, for example, has a disproportionately large proportion of the available space whereas Asia is running out of addresses. The remedy so far has been to use so-called Network Address Translations, or NATs, where a number of machines hide behind a gateway to the Internet. However, there are at least two problems with this. One is that the machines 'hiding' behind a NAT gateway do not have unique global addresses. A second, related problem is that use of NATs complicates end-to-end security. IPv6 will solve this by extending the address field to 128 bits, giving what is in practice an unlimited number of globally unique addresses.

Autoconfiguration

In IPv4, addresses need to be set up manually or using a protocol known as DHCP. By contrast, IPv6 allows autoconfiguration of the address, thus greatly simplifying network management. The self-generated IPv6 address combines a local network identifier and an identifier generated by the node in a 'plug-and-play' fashion. By simplifying the maintenance of large networks of simple devices, such as sensors and actuators, this automated procedure paves the way for the large-scale introduction of Internet technology into industrial environments. In fact, if the development 'IP all the way to the light bulb' is to become, as many predict, reality, this is an absolute requirement.

> The extensive IPv6 address range alone is enough to ensure its eventual introduction on a large scale in numerous applications.

Support for real-time traffic

The IPv6 header contains two fields specifically targeted at real-time applications, the 'Traffic Class' and the 'Flow Label'. The Traffic Class is an 8-bit field that enables routers to distinguish packets of different classes or priorities. The field is similar to IPv4's 'Type of Service' field, and one possibility would be to use it in combination with DiffServ (Differentiated Services). However, so far the field's exact use is not specified and research is still in progress to specify its operation. The Flow Label is a 20-bit label identifying flows. The IPv6 specification describes a flow as a sequence of

IPv6 – the whys and wherefores

IPv6, or Internet Protocol Version 6, was recommended by the Internet Engineering Task Force (IETF) meeting in Toronto in 1994. In 1998, the core set of IPv6 protocols became an IETF Draft Standard.

IPv6 is a new version of IP, designed to be an evolutionary step from IPv4, the version of the Internet Protocol currently in use. It can be installed as a normal software upgrade in Internet devices and is interoperable with IPv4. Users will be able to upgrade their hosts to IPv6, and network operators deploy IPv6 in routers, with a minimum of coordination between them.

IPv6 fixes several of the problems now being encountered with IPv4, such as the limited number of available addresses. And it adds improvements in areas such as routing and network configuration. IPv6 is expected to gradually replace IPv4, with the two coexisting for a number of years during a transition period.

IPv6 implementations are being developed for many different host operating systems and routers. Many common Internet applications already work with IPv6, and more are being ported all the time.

More information on IPv6 can be found at www.ipv6.org.

1	Respective headers of the IPv6 and IPv4. The address length of the IPv6 (bottom) has been augmented, while fields defined in IPv4 (top) that have turned out to be of little use have been removed.

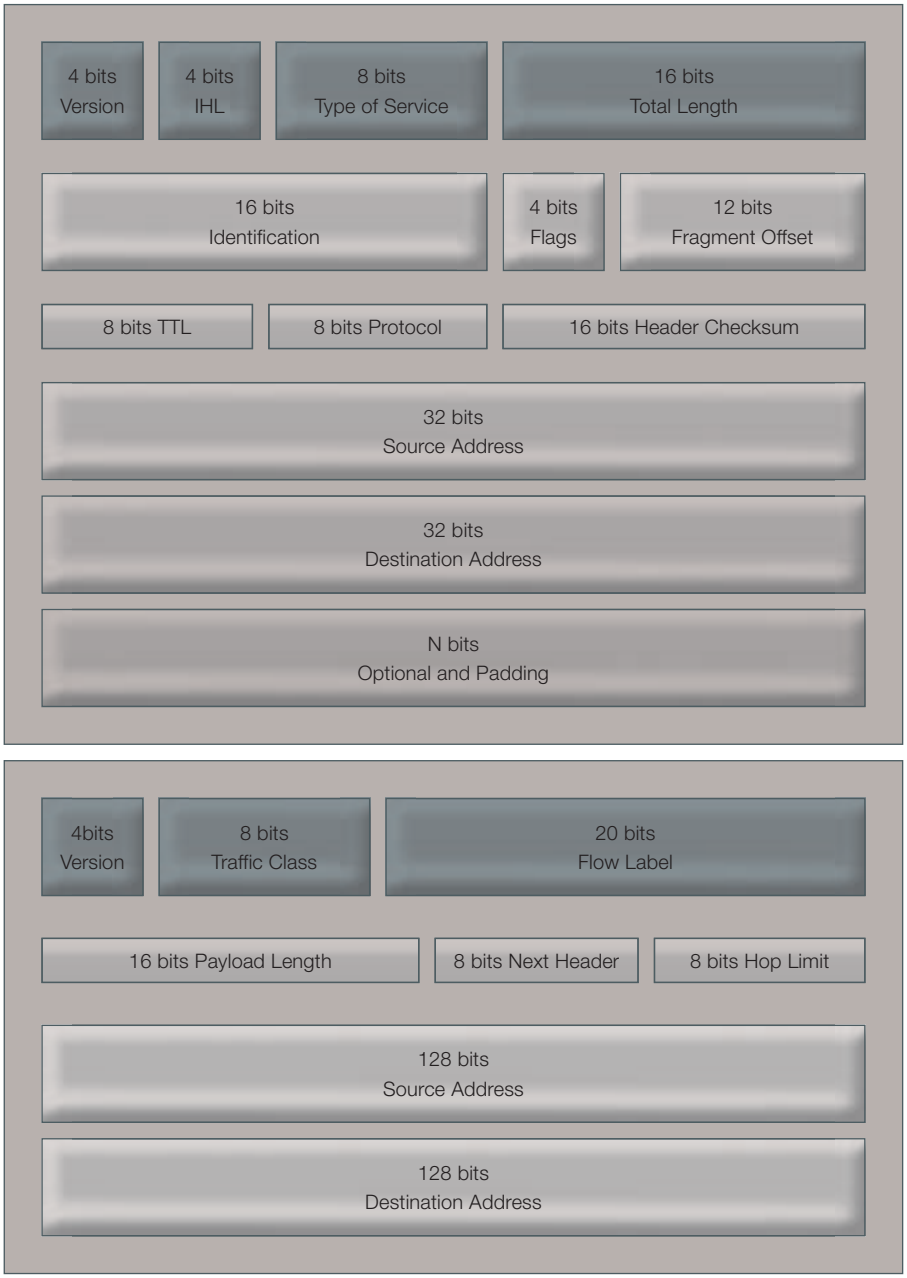

packets sent from a particular source to a particular (unicast or multicast) destination for which the source desires special handling. Special handling is here exemplified as non-default quality of service or real-time service. Flow labels can, for example, be used in combination with a resource reservation protocol such as RSVP. Again, the use of the Flow Label, although potentially potent, is still largely unclear.

Enhanced support for mobility

The framework defined for IPv4 provides mobility to nodes. The concept relies on the idea of each node having a home network that keeps track of its whereabouts. Whenever a node logs onto a foreign network, a software module known as the foreign agent (FA), provides a temporary address and informs the home network, or home agent (HA) of it. With IPv6, the foreign agent has been moved into the mobile node as an integral part of the protocol. This essentially enables all IPv6 enabled nodes to roam freely, regardless of the software installed in the visited networks. (It should be noted that full mobility is also possible with IPv4 using commercially available software supporting the mobility standard. Such a network has been installed at Corporate Research in Norway, where it provides internal mobility to selected nodes.)

Improved security

The security architecture for the Internet Protocol, known as IPSEC, provides the tools (encryption, authentication and key negotiation) needed to create virtual private networks (VPN) and let employees work from home (or any other location) without jeopardizing security. Whereas IPSEC is available as an extension to IPv4, it is part of the IPv6 protocol itself, ensuring its availability in all IPv6 nodes. Recently, however, certain concerns have been voiced about security in the mobility aspects of IPv6, MIPv6. A security flaw would essentially allow a malicious user to divert traffic for a node. This is an important issue and one that is currently under investigation by the standardizing committee.

Better provisions for multicasting

As online seminars, conferences, and similar applications become common the need for multicast support will increase if network bandwidth is to be preserved. In IPv4, multicast relies on a protocol known as the Internet Group Management Protocol (IGMP), which specifies its own message types. IPv6's equivalent of IGMP, called Multicast Listener Discovery (MLD), is based on ICMPv6, which must be implemented in every IPv6 node. This represents a major enhancement of multicast support compared with IPv4, which does give such support. Again, however, it has to be said that the basic functionality exists in IPv4 as an add-on, whereas in IPv6 it is an integral part of the protocol.

What are the implications for ABB?

A number of the key attributes of IPv6 are of obvious interest to ABB. Indeed,

as the Internet is increasingly introduced at or close to device level, aspects like autoconfiguration and multicasting will be of paramount importance. Also, given the large number of nodes that may conceivably be connected together in a system, its large address range provides some assurance that the threat of an address shortage could be a thing of the past. Finally, since the real-time requirements of industrial applications are often rigorous, the improved real-time support may turn out to be of interest.

To evaluate the protocol ABB first checked its performance in terms of delay through the stack. Given the improved header structure, a shorter delay should be expected. Then, in order to analyze the possible implications of IPv6 for ABB, we considered a typical scenario involving self-configuration of a wireless sensor network and large fieldbus networks with a relatively low bandwidth (such as EIB or LON).

Performance evaluation

The initial aim was to determine the performance of current IPv6 equipment. Two isolated test networks – one with the PCs run back-to-back and one where the packets pass through a router – were set up in the laboratory at the ABB research center in Norway 2. The PCs ran the Linux operating system and had a resident software implementing an IPv6 stack. The router was a Cisco 2514 with upgraded IOS to support basic IPv6 functionality. As these networks are extremely simple it is easy to identify the contribution of the different network elements to the observed delays. The round trip time (RTT) was measured for both IPv4 and IPv6, using the two network configurations shown, in order to compare the delay in the software, known as 'stack latency'.

2 Isolated test networks – one with the PCs run back-to-back (top) and one where the packets pass through a router (bottom). These extremely simple networks are used to determine the performance of current IPv6 equipment.

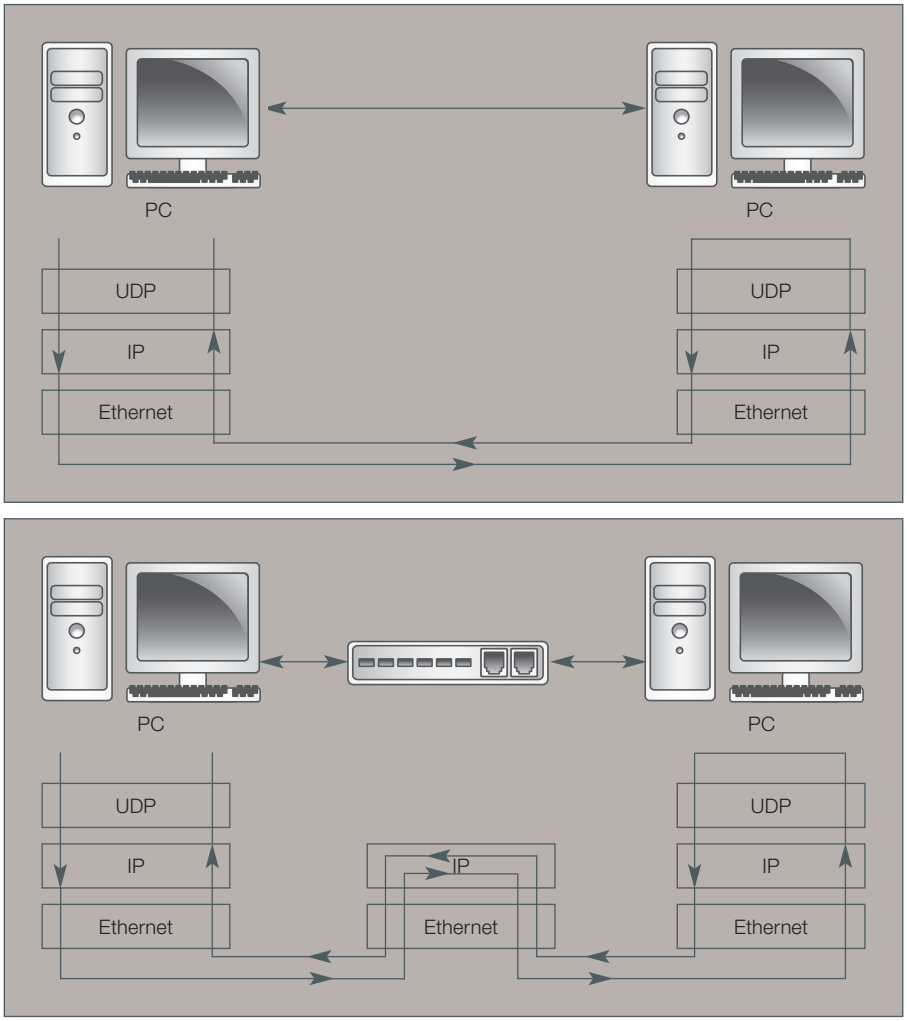

Several key attributes of IPv6 are of obvious interest to ABB. As the Internet is introduced at or close to device level, aspects like autoconfiguration and multicasting will be of paramount importance.

The results show that the two protocols perform roughly the same when the two PCs are run back-to-back 3, with IPv4 exhibiting slightly lower latency than IPv6. Both stacks are software implementations, and the difference in performance may be explained by the fact that the IPv6 solution is a recent product, whereas the IPv4 stack has undergone a number of performance enhancements over the years. When the packets pass through the router, however, the measured RTT of IPv6 is significantly longer than for IPv4. This difference is caused by the fact that the router uses optimized hardware to forward IPv4 packets, while IPv6 forwarding is still done in software.

The only conclusion that may be drawn from these observations at this point is that current versions of IPv6 are sub-optimal and do not reflect the full potential of the protocol. IPv6 is a high-performance protocol, with a header structure designed for rapid processing, but the current software solution that was used does not perform as well as its IPv4 hardware counterpart.

Wireless and wired sensors and networks

Consider the case of a factory floor with an assembly line on which new items are continually entering and leaving. Sensors or tags attached to moving parts may have to be given an IP address as they enter the coverage area of the base stations. In such a situation the autoconfiguration properties of IPv6 could be interesting as they allow new network elements to automatically take a new address that is guaranteed to be unique. In IPv4, however, the traditional approach is to use DHCP. In a standard due to be released soon, provisions are made for automatic generation of link local addresses in IPv4. In this scenario each new node would, however, have to check for possible address collision with every other node in the network. All the nodes would then have to be woken in order to run a collision detection check, thereby wasting valuable battery power on the sensors. A benefit of an IPv6 implementation could therefore be its potential to reduce the power at the sensors, which could stay dormant.

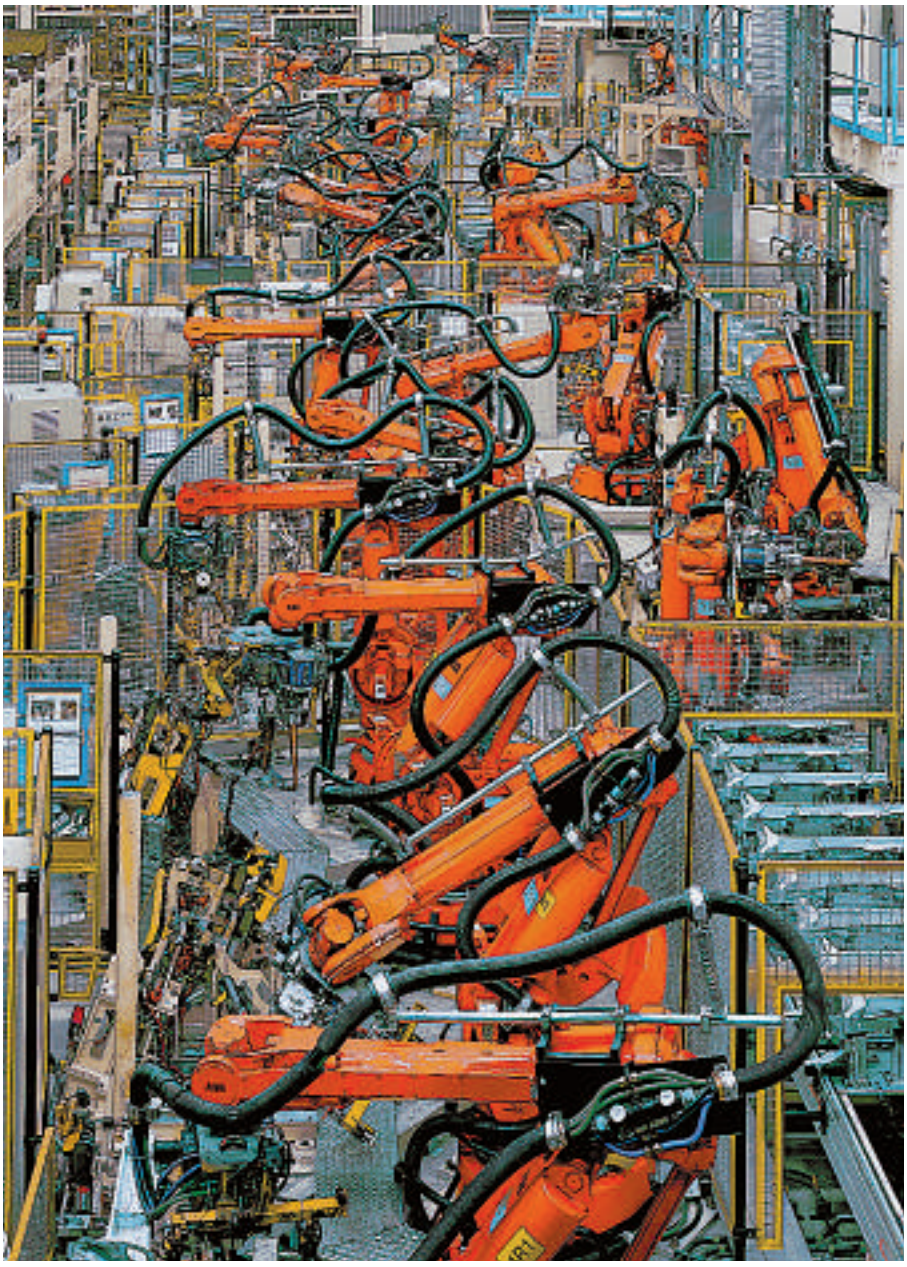

Now let us suppose that the wireless network is implemented as a Bluetooth [2] piconet or scatternet. Assuming no retransmissions and no collisions, it can be determined that the average power consumption in each sensor due to address collision detection is a linear function of the entry rate. It has also been shown that the additional power consumed due to address collision detection is relatively modest. Thus, in this particular instance the gain with IPv6 is of limited value.

In the wired part of the network the number of nodes can easily extend to several thousand, as in building applications. These nodes are typically, although not necessarily, cheap devices with little processing power and only limited memory. For such devices the issues of primary importance are computational complexity and memory requirements.

As we have already seen, IPv6 has quite a few mandatory features that are optional in IPv4 (such as support for mobility and security.) The inclusion of these features will almost certainly lead to a stack with a larger footprint, which is exactly what we want to avoid. On the other hand, the simplified header structure will eventually reduce the load on the processor.

Another important issue in low bandwidth networks is the size of the packets to be transmitted. IPv6 headers are 40 bytes compared with 20 bytes for IPv4. For large capacity networks this small additional overhead is not dramatic. But for bandwidth limited systems such as EIB (running at 9.6 kbps) with small payloads, such an overhead is significant and systems quickly become saturated.

If mobility, enhanced security, or multicasting is required between the nodes there is always the possibility of using these parts of the IPv4 specification. IPv6 is not strictly needed (although possibly useful) as all the network elements are under our control. The question is, of course: does having these options outweigh the inconvenience?

This is an open question that will depend on the application, but under normal circumstances it is doubtful. Real-time requirements of nodes may sometimes be rather strict, in the order of 20 ms. The improved focus on real time in IPv6 could prove helpful, but as indicated in the previous section, the exact use of this capability is still largely unresolved. It could well turn out that it becomes one of IPv6's major assets, but it is too early to say. Only later specifications (and actual use) will decide if this is indeed the case.

For users faced with a large number of nodes, autoconfiguration is a must and this feature of IPv6 would definitely be appreciated. As stated above, however, a specification for IPv4 will solve autoconfiguration with this protocol, albeit with some network traffic overhead.

The increased address range is the one 'killer argument' for introducing IPv6. If our nodes are to have globally unique addresses only a few years separate us from the time when there will be no choice; there simply will not be enough addresses left.

In short, in the example we have looked at there is no compelling argument for introducing IPv6 other than possibly the enhanced address range and real-time performance. Even these arguments appear uncertain: the network can usually hide behind a NAT gateway and the handling of real-time sources is uncertain.

And the conclusions ABB has drawn?

ABB's study of the implications of IPv6 shows its performance in terms of latency to be inferior to that of IPv4, but this is expected to improve as the implementations mature. Most of the advantages of IPv6 are functionality that may also be achieved with IPv4 in systems where we have complete control of the nodes. This is often the case in factory automation applications. In more heterogeneous networks the case for IPv6 is much stronger. The single most important aspect of IPv6 is the extensive address range. This alone is enough to ensure its eventual introduction on a large scale in numerous applications.

3 Probability (p) of round trip times (in µs, horizontal axis), measured for IPv4 (black) and IPv6 (blue) using the two network configurations in figure 2.

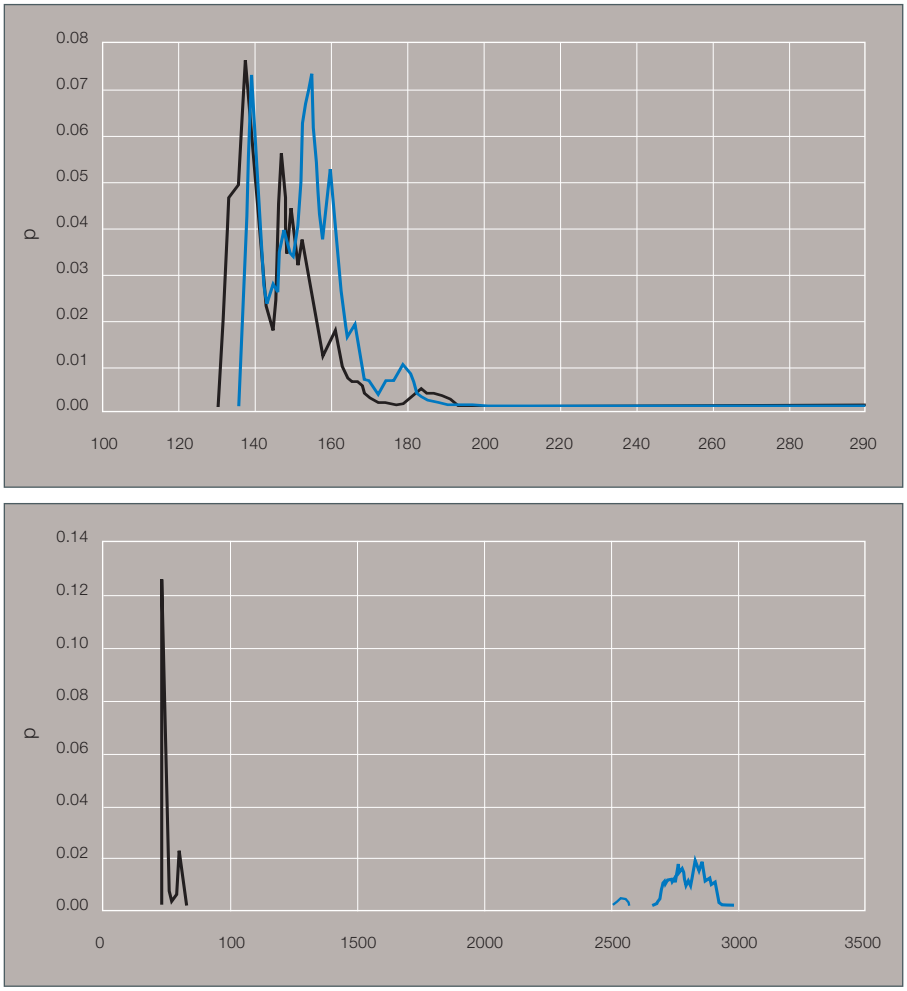

The top graph shows roughly the same results with the two PCs running back-to-back. The largely different results for the configuration with router (below) are due to optimized hardware being used to forward IPv4 packets, while IPv6 forwarding is still done in software.

Niels D. Aakvaag
ABB AS
Norway
niels.aakvaag@no.abb.com

Nils A. Nordbotten
University of Oslo
Norway
nilsno@ifi.uio.no

[1] S. Deering, R. Hinden: Internet Protocol Version 6 Specification. RFC 2460, December 1998.
[2] Bluetooth specification, version 1.1. More information on Bluetooth can be found at www.bluetooth.com.