# Porting TCP/IP to the 802.4 Token Bus LAN.

*Andrew McRae*

Megadata Pty Ltd.
*andrew@megadata.mega.oz.au*

## The 802.4 Token Bus LAN.

The OSI standards embrace three Local Area Network media specifications, each of which cater for different environments and user needs. Each specification details the physical and electrical requirements of the LAN, as well as the method of accessing the LAN, known as the Media Access Control (MAC).

The IEEE 802.3 specification (Ethernet to all intents and purposes) formalised an existing LAN defacto standard developed some time ago, which (understating somewhat) has become widespread in its installed base and vendor support. Ethernet is classed as a Carrier Sense Multiple Access/Collision Detect (CSMA/CD) technology. Ethernet operates at 10 Megabits/second, and due to the nature of the MAC, has limits that should be adhered to for correct operation (though these limits are often abused).

The 802.5 Token Ring network is also designed for the office environment, and has achieved a limited success, mostly due to the support of certain large vendors. It is cabled as a physical ring, and can operate at 4 Megabits/second with later versions operating at 16 Megabits/second. The stations are connected to the ring via active transceivers. The MAC is based on a token being transmitted from each station to the next. Stations can only transmit when they ‘own’ the token.

Both Token Ring and Ethernet are well known, and have a large installed base. The 802.4 Token Bus LAN standard is relatively unheard of in the computing community, except perhaps in the process control industry. The 802.4 Token Bus network is designed as a physically and electrically robust LAN for running the Manufacturing Automation Protocol (MAP), and is aimed at providing LAN connectivity at the ‘factory floor’ level. Historically, General Motors was the prime factor behind MAP, as a way of standardising the interconnection of process control equipment, numerically controlled machining equipment and factory robots throughout a plant environment.

MAP itself is defined as a complete OSI protocol stack, and due to the considerable size and complexity of MAP, and the need to fit working MAP implementations into factory equipment, a revised version of MAP (called *Mini-MAP*) was defined that removed certain protocol layers.

Token Bus is a hybrid of LAN technologies. A range of different media, speed and cabling options are available, allowing selection of the most appropriate technology for various needs. Token Bus can operate on 10 Megabit/second broadband, 10 Megabit/second carrier band, 5 Megabit/second carrier band, and 1 Megabit/second shielded twisted pair. The physical cabling is arranged as a bus, allowing easy extension or partitioning (Figure 1). The cable used is armoured, and has extremely good electrical tolerance to noise and interference.

Each node is connected to the LAN via a passive tap. The tap is hermetically sealed and is impervious to dust and vibration. Since the tap contains no active components, much better electrical isolation can be obtained, as well as protection from LAN disintegration due to tap electrical failure.

Whilst the physical LAN is a bus, the MAC is achieved through a logical station ring, where a token is passed from station to station (Figure 2). Since the LAN is a bus, the media is a broadcast medium, but with a layer of sophisticated control dealing with the access of the stations onto the LAN via a logical token ring. This allows greater control over the operating characteristics of the LAN, both dynamic and static, as well as exceptional robustness and fast response to error conditions.

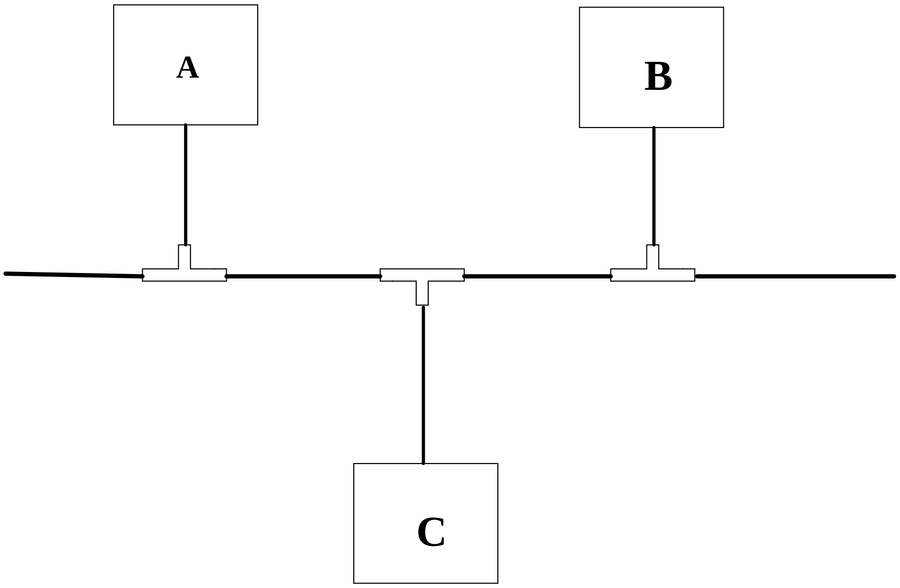


**Figure 1**

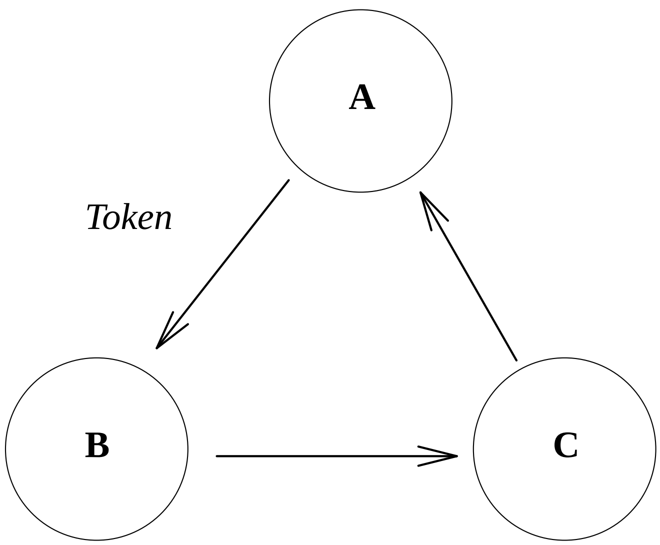


**Figure 2**

Token Bus demands quite a bit more from a controller than Ethernet. For example, the Token Bus controller after passing a token to the next logical station will eavesdrop on the LAN to ensure that the token pass was successful. If not, it retries and eventually recovers by cutting out the failed station from the ring and establishing a new ‘next station’ by passing the token to the next logical station after the failed station. Periodically a delay is inserted in the token pass to solicit for any new stations that wish to enter the ring.

The ring is established upon initialisation by stations attempting to ‘claim’ the token. The 48 bit MAC address is instrumental in providing prioritisation of the stations both with token claim and also establishing the order of the token pass (from highest MAC address down to lowest MAC address). Major errors such as a station detecting another station using the same MAC address will force the offending stations into an OFFLINE state so they are not part of the ring. Stations can only transmit when they hold the token. If any station detects the collapse of the logical ring (by not hearing any token passing), that station will attempt to reinitialise the ring. The ring set up time is short enough (depending on the number of stations it can be as low as a small number of milliseconds) so that LAN access is not greatly affected in the event of a collapse and reinitialise.

Apart from the robust nature of the physical bus and the logical ring, the Token Bus LAN allows *deterministic* control over the LAN access for each station so that certain and definable minimum times can be guaranteed. A widely accepted (and in some cases valid) criticism of Ethernet is the poor reponse under load, and the inability to prioritise access to the LAN. Token Bus as a deterministic LAN guarantees that even in the event of every station wishing to transmit lots of data, no one station will hog the LAN, and every station will be given a chance to send; moreover, the determinism is under the control of the system designer.

This attribute is a key feature of 802.4, and is a major factor in the choice of 802.4 as the preferred LAN for real time or mission critical distributed systems. How is this LAN determinism achieved?

The LAN is configured for a specified Target Rotation Time (TRT) in octets; this is the desired maximum time it takes for a token to complete a full cycle of the logical ring, expressed as the number of octets transmitted on the LAN, including all preamble and frame control octets. When a station receives the token, it considers the number of octets that have been transmitted since it received the token *last* time. This value (termed the Last Token Rotation Time - LTRT) is compared against the TRT and determines whether this station can transmit any packets. If a station has no data to transmit, it simply passes the token, so the minimum token rotation time may be much less than the TRT maximum. However if a station has transmit packets queued it will only begin the transmission of each if the number of octets that has been transmitted on the LAN is less than the desired TRT. In other words, every station monitors the traffic volume, and automatically restricts its LAN usage in the event of congestion. The net effect is an averaging of available bandwidth to all stations fairly.

As an example, say we have 5 stations, called A, B, C, D and E. The token is passed from A→B→C→D→E. The LAN has a maximum TRT of 2000 octets (not including framing and preamble octets), and each station is transmitting the following packets every rotation of the token:

| Station | Packet size | count |
|---|---|---|
| A | 500 | 1 |
| B | 300 | 2 |
| C | 100 | 2 |
| D | 200 | 1 |
| E | - | none |

The table below details each station's last measured rotation time, and the number of data octets it transmits on each rotation.

| | | A | B | C | D | E |
|---|---|---|---|---|---|---|
| 1 | LTRT | - | - | - | - | - |
| | XMIT | 500 | 600 | 200 | 200 | 0 |
| 2 | LTRT | 1500 | 1500 | 1500 | 1500 | 1500 |
| | XMIT | 500 | 600 | 200 | 200 | 0 |
| 3 | LTRT | 1500 | 1500 | 1500 | 1500 | 1500 |
| | XMIT | 500 | 600 | 200 | 200 | 700 |
| 4 | LTRT | 2200 | 1700 | 1400 | 1600 | 1400 |
| | XMIT | 0 | 300 | 200 | 200 | 700 |
| 5 | LTRT | 1400 | 1900 | 1900 | 1800 | 2000 |
| | XMIT | 500 | 300 | 100 | 200 | 0 |
| 6 | LTRT | 1100 | 1100 | 1400 | 1500 | 1500 |
| | XMIT | 500 | 600 | 200 | 200 | 700 |

The first two rotations are less than the maximum TRT, so all stations transmit the data they have queued without delay. On the third rotation station E has a 700 byte packet queued for transmission; because the measured rotation time is less than the TRT, it sends the packet. However station A now has a measured rotation time of 2200 octets, and so passes the token without sending any data. Station B has a measured time of 1700 octets, and so can only transmit one of its two 300 byte packets. Station A gets to transmit on the next rotation though, whilst station E will miss out, and so on. In effect the transmitting of the data will average so that the TRT is maintained, but no station is locked out of sending data for more than one rotation.

The MAC also defines a 4 level priority scheme so that a much finer degree of control can be placed over the basic Token Bus access. Token Bus is becoming increasingly common over slower speed networks such as 1 Megabit twisted pair, and even down to 100 Kilobit twisted pair. At these speeds, the priority mechanism is absolutely vital so that maximum control is gained over the limited bandwidth resource. It guarantees that more important data are sent ahead of less important data, especially at times when there is a burst of LAN activity (which is often the time when the important data has to get through fastest).

**Some Results.**

The Token Bus network was implemented using Megadata Distributed Interface Units (DIU). The experimental network was set up as shown in figure 3. A Serial Line/Internet Protocol (SL/IP) connection was used to gateway an Ethernet network containing a variety of UNIX† workstations such as Sun Sparcstations, Hewlett Packard workstations, PCs etc. An Annex terminal server acted as the IP router for forwarding network packets.

The use of TCP/IP to interconnect the DIU via SL/IP to a an Ethernet meant that all the network debugging tools such as *ping*, *ttcp* and others could be used from any workstation on the network. One of the first applications to be developed was a remote login facility via *telnet*, which was an excellent proof of concept test. Users could remotely log onto the DIU, and perform system configuration, examine statistics and monitor operation, all from the comfort and safety of their workstation. It also allowed easy prototyping of application software, as well as fast downloading of programs and data.

Further measurements indicated that a basic Remote Procedure Call (RPC) between two DIU's took between 5 and 10 milliseconds, which was pleasantly surprising given the capabilities of the CPU.

Prospect Electricity in Sydney's west has worked in conjunction with Megadata to formulate and specify a substation automation and control system, that implements a complete environment for automatically controlling and managing an electrical distribution substation. Megadata DIUs are physically distributed throughout the substation, on or next to the actual equipment being controlled, and are networked via TCP/IP and Token Bus. Each DIU contains control logic for the type of equipment it is connected to, and all nodes share data via a distributed database. This is currently undergoing final test, and is due to be

† UNIX is a registered trademark of AT&T UNIX System Laboratories.

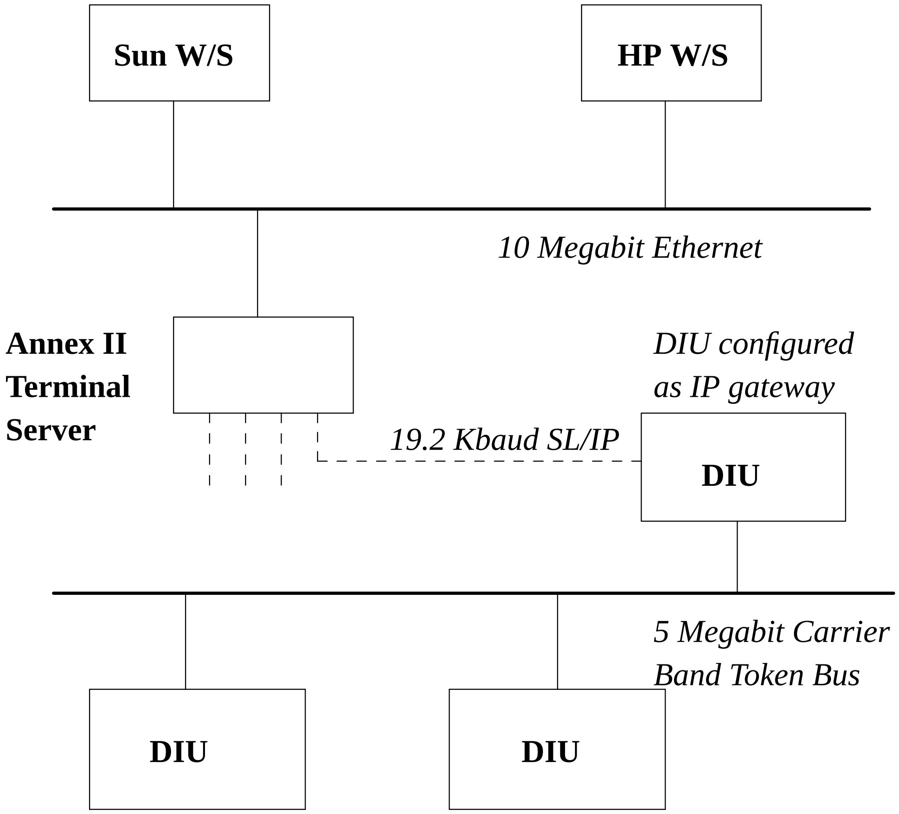


**Figure 3**

installed at sites throughout Western Sydney, and in tests has shown to be capable of high speed control, data acquisition and monitoring, and yet has sophisticated facilities for connection to a Wide Area Network via TCP/IP, remote login to any DIU from any other node, and can be connected to any equipment that supports TCP/IP. The performance is entirely due to the distributed nature of the system, and to the fact that TCP/IP is a fast and lightweight protocol stack.

In such a system it is possible to build an enterprise wide network based on Open Systems protocols which interconnects the largest mainframes with the smallest embedded nodes.

**Future Applications.**

A number of future applications are being investigated at the present time.

It is envisaged that the IP priority options can be directly mapped to the 802.4 priority scheme, so that the application can specify the priority of the data being sent. This would allow a great degree of control over the guaranteed delivery times of network packets.

To further enhance the networking application environment, a number of extra network facilities are being ported, such as Sun Microsystems's RPC/XDR remote procedure call system, providing applications with a sophisticated and powerful RPC mechanism, and also an SNMP agent conformant with the standard MIB for 802.4 Token Bus interfaces, so that each node can be managed using standard network management tools.

The DIUs may be connected directly to Ethernet, which allows a much faster gateway between 802.4 and Ethernet to operate. Developments are being considered which provide for more powerful CPUs and memory expansion, so that sophisticated applications beyond the capabilities of the current product can be run. Other possibilities include employing current generation embedded processors to provide a smaller, cheaper node, but offering powerful CPU performance and floating point in a single chip.

The logical outcome is that Token Bus networks can be internetted to other networks, and IP packets routed across the gateway rather than slow application bridges performing data conversions. Thus it would be possible to closely link the UNIX workstation world with the factory floor environment, facilitating the use of a wide range of technology and computing power to access the real time applications, was well as delivering plant data to MIS databases. The factory floor processing nodes can be part of a complete asset management and control system, that allows the entire organisation access to real time data, without limitations caused by the use of proprietary equipment or protocols. The provision of more powerful processors and significant amounts of memory allow the integration of TCP/IP and MAP as a dual protocol stack architecture. Application and transport bridges would then allow many different types of systems to inter-operate and communicate via the dual protocol product.

Organisations are now starting to realise that data acquisition and control is not a separate problem that demands a separate solution (as in the past), but that real time data needs to be *integrated* into the entire enterprise's data model, so that effective planning and management can take place.

Examples of such systems are Substation automation and control, water treatment plant management, power station managament, factory floor and automated plant environments, and pumping stations.

The ultimate goal of employing Open Systems standards is to allow users to preserve their investment (hardware and applications) by easily and simply interconnecting computing equipment and software from a range of vendors. Whilst this is starting to become a reality in the workstation and general computing fields, it is unheard of in the the real time field. What should be happening is for organisations to treat real time acquisition and control as similar problems to the general data processing problem (in terms of Open Systems connectivity and standards), albeit with specific criteria applied such as availability and response time. In a similar fashion, vendors of control and plant equipment should be *building in* Open Systems connectivity into products e.g. when a manufacturer sells a circuit breaker or a large turbine, it comes with a standard real time LAN (e.g. Token Bus) connection, allowing the installer to simply plug the equipment onto the data acquisition LAN and immediately integrate it into the control network.

**Conclusion.**

The use of TCP/IP as an Open System protocol suite on the IEEE 802.4 Token Bus LAN in the real time embedded control area has proved to be a success, displacing proprietary protocols. The use of a standard LAN means that, in time to come, the physical and logical interconnection of different vendor's equipment will be easier. The porting of TCP/IP to 802.4 provides a fast and efficient networking standard which can operate on a limited hardware platform, whilst still maintaining the numerous advantages of a common networking foundation across a complete range of computing environments.