Fibre-Ribbon Pipeline Ring Network with Distributed Global Deadline Scheduling

Carl Bergenhem[1], Magnus Jonsson[1], and Jörgen Olsson[2]

1. {carl.bergenhem, magnus.jonsson}@ide.hh.se, Computers and Communications lab, Halmstad University, Halmstad, Sweden. Phone: +46-35-167100 Fax: +46-35-120348
2. jool@vtd.volvo.se, Volvo Technological Development Corporation, Göteborg, Sweden. Phone: +46-31-7724681

Abstract

This paper introduces a novel, fair medium access protocol for a pipelined optical ring network. The protocol provides global optimisation of deadline constraints on a packet basis. Requests for sending packets are sent by the nodes in the network to the current master node. The master uses deadline information in the requests to determine which request is most urgent. Arbitration is done in two steps, collection and distribution phases. The protocol is therefore called two-cycle medium access (TCMA). The network is best suited for LANs and SANs (system area networks) such as a high speed network in a cluster of computers or interconnection networks in embedded parallel computers. Services possible in this network include best effort messages, real-time virtual channels, functions used in parallel processing such as barrier-synchronisation, and functions for reliable transmission. These are possible without additional higher level protocol layers. A simulation analysis of the network with the novel protocol is presented. Further analysis shows worst case latency, minimum slot length, and fairness of the protocol.

Keywords: Real-time communications, protocol, global deadline optimisation, packet constraints, fibre-optic interconnection network, embedded systems, parallel processing

1 Introduction

The contribution put forward by this paper is a novel medium access protocol that uses the deadline information of individual packets, queued for sending in each node, to make decisions, in a master node, about who gets to send. The new protocol may be used with a previously presented network topology; the control channel based fibre ribbon pipeline ring (CC-FPR) network [1]. Simulations of the protocol prove its validity.

The proposed medium access protocol provides the user with a service for sending best effort messages for which the timing constraints are globally optimised. Because of this property the protocol is suitable for real time communication. Further more, the global optimisation is a mechanism that is built into the network protocol. No further software in upper layers is required for this service. The scheme of globally optimising deadline constraints presents an advantage over networks that don't arbitrate medium access on this property. Some examples of this may be found in [1], [2], [3], [4], [5]. However, these networks may have upper layer protocol added to them to give them better characteristics for real-time traffic but this increases the complexity of the system. Also, it is hard to get fine deadline granularity using upper layer protocols.

Real-time services in the form of best effort messages, as mentioned above and real-time virtual channels (RTVC) are supported for single destination, multicast and broadcast transmission by the network. A service for slot reservation, used for hard real-time traffic such as RTVCs is also provided [1]. The network also provides services for parallel and distributed computer systems such as short messages, barrier synchronisation and global reduction. Reliable transmission service (flow control and packet acknowledgement) is provided as an intrinsic part of the network [6], [7].

The network with the proposed protocol is best suited for LANs and SANs (system area networks) where the number of nodes and network length is relatively small. This is important since the propagation delay adversely affects the medium access protocol. An example of a suitable application is in an embedded system, e.g., for use as an interconnection network in a radar signal processing system, or as a network for use in cluster parallel computing. A problem with the original CC-FPR protocol [1] is that a node does not consider the time constraints of packets that are queued in downstream nodes. The novel network presented here does not suffer from this problem.

Novel optical components result in the possibility of new network solutions for the increasing bit rate demands of parallel and distributed systems. Motorola OPTOBUS™ bi-directional links with ten fibres per direction are used but the links are arranged in an unidirectional ring architecture where only $\lceil N / 2 \rceil$ bi-directional links are needed to close a ring of N nodes (assuming that N is an even number). Fibre-ribbon links offering an aggregated bit rate of several Gbits/s have reached the market [8]. The increasingly good price/performance ratio for fibre-ribbon links indicates a great success potential for the proposed type of networks.

The physical ring network is divided into three rings or channels (see Figure 1). For each fibre ribbon link, eight fibres carry data, one fibre is used to clock the data, byte by byte, and one is used for the control channel. Access is divided into slots like in an ordinary TDMA (Time Division Multiple Access) network. The control channel fibre is dedicated for bit-serial transmission of control-packets, which are used for the arbitration of data transmission in each slot. The clock signal on the dedicated clock fibre, which is used to clock data, also clocks each bit in the control-packets. Separating clock- and control-fibres simplifies the transceiver hardware implementation [9]. The control-channel is also used for the implementation of low-level support for barrier-synchronisation, global reduction, and reliable transmission [6].

The ring can dynamically (for each slot) be partitioned into segments to obtain a pipeline optical ring network [2] where several transmissions can be performed simultaneously through spatial bandwidth reuse, thus achieving an aggregated throughput higher than the single-link bit rate (see Figure 2 for an example). Even simultaneous multicast transmissions are possible providing multicast segments do not overlap.

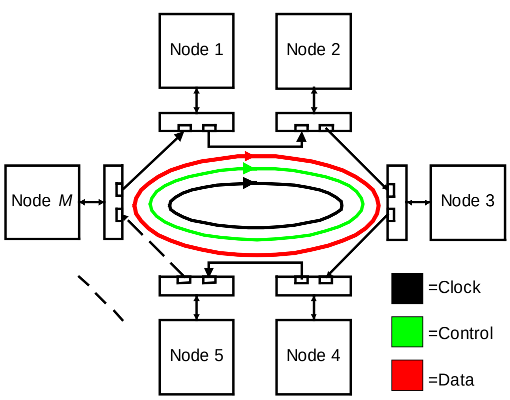

Figure 1: A Control Channel based Fiber Ribbon Pipeline Ring network.

The rest of the paper is organised as follows: Section two briefly describes the original protocol since many ideas are similar to the presented. Section three presents the novel medium access protocol. Section four gives a short description of RTVCs and how they are implemented. Section five presents an analysis of worst case latency and minimum slot time, and describes the simulations carried out. These results are presented and discussed. Conclusions are presented in section six.

2 Original protocol

The novel medium access protocol presented and analysed in the following sections, use many ideas from the CC-FPR protocol presented in [1]. Therefore some details of the CC-FPR protocol are presented here.

The original CC-FPR network protocol is insensitive to propagation delay in the sense that no feedback, from nodes back to the master, is needed during arbitration of the network. Arbitration is decentralised by having nodes take equal, round-robin, turns to be master. The protocol is based on the use of a control-packet that, for each slot, travels almost one lap (over N–1 links) round the control-channel ring. In the time domain the control-packet always travels around the ring in the time-slot preceding the time-slot for which it controls the arbitration. The control-packet will hence always pass each node one time-slot before the data-packet it is related to passes.

At the beginning of each slot the master initiates a control packet that contain its own needs for packet transmission. Each node succeeding the master checks the control-packet when it passes the node to see: (*i*) if it will receive a data-packet in the next slot, and (*ii*) if a data-packet will pass the node in the next slot. If no data-packet will pass the node, i.e., the rest of the ring back to the master is free, the node will have the possibility to

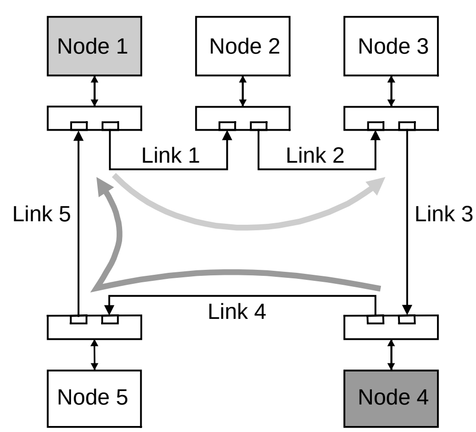

Figure 2: Example where Node 1 sends a single-destination packet to Node 3, and Node 4 sends a multicast packet to Node 5 and Node 1.

transmit a data-packet in the next slot in this segment of the ring. The node signals this by changing the control packet to reflect how it will send. Because all the nodes succeeding the master repeat the procedure of checking the control-packet for the possibility to send, multiple transmissions in different segments of the ring might be possible in the same slot. Observe that the control packet does not return to the master and that a data packet cannot be transmitted across the master node since the clock is interrupted there.

A problem with the original CC-FPR protocol is that a node does not consider the time constraints of packets that are queued in downstream nodes. See Figure 2 for an example described below. Node one decides that it will send and books links one and two, regardless of what Node two may have to send. This means that packets with very tight deadlines may miss their deadlines. For a further presentation of the original protocol, refer to [1].

3 Two-cycle medium access protocol

The greatest difference in function between the new and the old protocol of accessing the network is that arbitration is done in two phases instead of one. The two phases to medium access are collection phase and distribution phase (see Figure 3). Therefore it is referred to as the two-cycle medium access protocol, (TCMA protocol). As can be seen, the protocol is time division multiplexed to share access between nodes. The basic time unit is called a slot and the minimum size of the slot is analysed in section five.

As in the original medium access protocol, the role of network master is cycled equally, round robin, around the ring. Thus all nodes are identical. The role as master is passed on to the next down stream node at the end of the slot. Every node detects when the clock signal is interrupted at the end of the slot and nodes have a counter to determine who is next master.

There are two types of TCMA control packets, which are used in each of the two phases (see Figure 4). A complete collection phase packet will contain a start bit and total of *N*-1 requests that are added one by one by each node. The master receives it's own request internally. Each request consists of three fields. The "prio"-field contains the priority level of the request. It is further described below. Nodes use the link reservation and destination fields to indicate destination node(s) and which links must be traversed to reach the destination node. Since a node may write several destination nodes into the destination field, transmissions may be multicast or broadcast. In the distribution phase packet the "result of requests"-field contains the outcome of each nodes request. This is the only field, in this phase, which contains network arbitration information. The others are used for services such as, e.g., reliable transmission ("ACK/NACK"- and "flow control"-fields) and, e.g., global reduction and short messages (the "Extra information"-field).

The time until deadline (referred to as laxity) of a packet is mapped, with a certain function, to be expressed within the four-bit limitation of the current version of TCMAs priority field. A lower priority implies a shorter laxity and thus a more urgent message. The priority field is a central mechanism of the TCMA protocol. The result of the mapping is written to the priority field (see Figure 4). A wider field of bits would provide higher resolution of priority and would probably have an advantageous affect on performance. Further

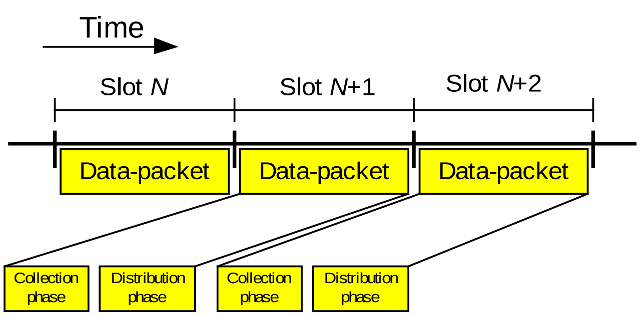

Figure 3: The two phases, collection and distribution, of the TCMA protocol. Notice that the network arbitration information, for data in slot *N+1,* is sent in the previous slot, slot *N.* Observe that the lengths of the phases, and placement in time, in the diagram are not to scale.

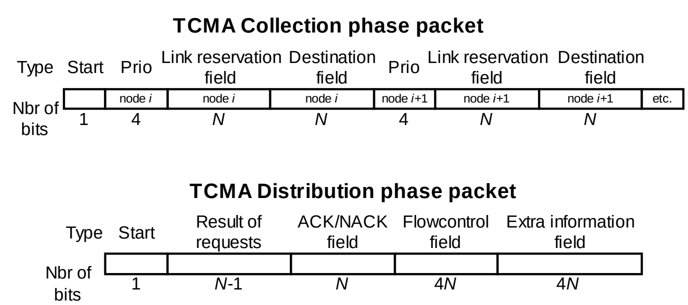

Figure. 4: Contents of the TCM control packets. Notice the possibility in the distribution phase packet to send information other than for medium access arbitration.

evaluation is out of the scope of this paper. Two mappings between deadline and priority, logarithmic and linear (see Figure 5), have been simulated. Results show a negligible difference in performance of throughput, packet-loss, and latency. Further evaluation of how the performance is affected by different mappings is therefore put beyond the scope of this paper. For the simulations presented in section five, logarithmic mapping is used. This mapping gives higher resolution of laxity, the closer to its deadline a packet gets.

All nodes including the master, have information about which slots have been reserved, e.g., for RTVCs between pairs of nodes. This is because a request to establish an RTVC is broadcast to all nodes. Therefore a node will only send a request that may be possible to fulfil regarding RTVCs in the own or other nodes that would use links in the path of the packet that the node would want to send. Observe that a node will not request a transmission that would be "across" the master since the clock signal is interrupted there. This implies that a request will only be rejected if requests from other nodes are more urgent. The node selects its most urgent message as the request. In the case that there are several messages that are equally urgent, the message that is destined furthest and possible to transmit in the next slot is selected. Slots belonging to RTVCs do not need to be "requested" since they are already reserved (see section four).

In the collection phase the node that is currently master generates an empty packet and transmits it on the control channel. Each node adds its request, in turn, for sending in the next slot. A request packet that will be added, by a node, to the packet from the master can be generated and ready to be sent as soon as the node knows the status of its queues so that only a very small delay is

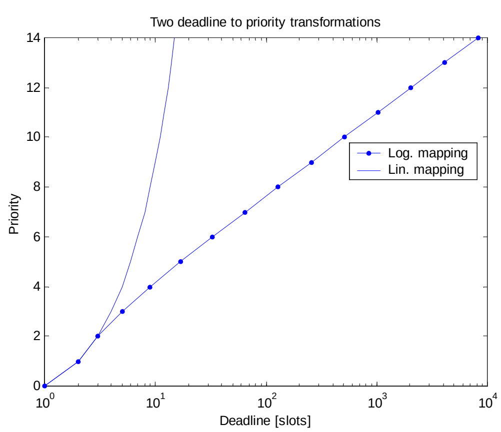

Figure 5: Two different deadline-to-priority mappings that were tested. For the linear transformation, deadlines longer than 14 slots are all mapped to priority level 14.

incurred. The additional delay is caused when the collection phase packet passes a node. It is assumed to be one bit time. The same delay also applies to the distribution phase packet. The request contains the following information: priority, link-reservation and destination field (see Figure 4). If the node has nothing to send, it signals this to the master by using a reserved priority level (15 in the proposed protocol) and zeros in the other fields.

When the completed collection phase packet arrives back at the master, the requests (including one request from the master) are processed. There can only be *N* requests in the master, as each node gets to send one request per slot. The list of requests is sorted primarily by priority and secondly by the distance to the destination, i.e., furthest within segment (FWS)[2]. The master traverses the list, starting with the request with lowest priority (closest to deadline) and then tries to fulfil as many of the *N* requests as possible. In case of priority ties, the request with the largest distance to its destination is chosen. If there still is a tie then the master's request has priority over other request. Because of spatial slot reuse, several requests may be granted permission to transmit in different segments, during the same slot, providing that segments do not overlap. This is also called pipelining of packets (see Figure 2).

When the master has scheduled the requests it distributes the result to all nodes in the distribution phase. In this phase the master node, and only the master node, has possibility to make use of the "extra" fields in the distribution phase packet such as sending acknowledges for packets sent during previous slots. For further explanation of this, see [6]. When all nodes have received the results of the request, each node is ready for the beginning of the next slot where data may be transmitted. A request was granted if the nodes "request result field"-bit in the distribution phase packet contains a "1". As in the collection phase, the distribution phase packet is delayed when passing each node. Again, this delay is assumed to be one bit time, although not as important as for collection phase since the master does not get feedback from nodes in the distribution phase. The master receives the result of the requests internally.

The advantage of this protocol is that the deadline requirements of all packets are taken into account and considered at a global level. Since the packets deadline information is collected into a "global queue" in the master, packets from each node can be sent in an earliest deadline first (EDF) fashion and the timing constraints of the packets can be seen as globally optimised.

4 Real-time virtual channels

Logical connections with guaranteed bit rate and bounded latency can be realised in the network by using slot reserving. Such connections are referred to as RTVCs. Either the whole ring is reserved for a specific node in a slot, or one or more segments of the ring are dedicated to some specific node(s). Slots are organized into cycles with a set number of slots. Nodes keep track of the current slot index in the cycles. A slot that has been reserved for an RTVC, guarantees that transmission is possible every cycle thus guaranteed bit rate. Several slots may be reserved for an RTVC in order to increase the guaranteed bit rate. Initially, each node has *J* non-reserved slots where it is master, giving a cycle length of $N \cdot J$ slots. So as to always have bandwidth for best effort traffic controlled by the TCMA protocol, described in section three, only *J*-1 slots are reserveable.

When a node wants to reserve a slot for an RTVC, it searches for slots where the required links are free, so allocation of a new segment can be done. First, the node's own slots are searched. If not enough slots could be allocated for the reservation, the search is continued in other nodes. In this case, the node broadcasts a packet containing a request to all other nodes to allocate the desired segment in their slots. The packet contains information about the links required and the amount of slots needed. Each node then checks if any of its own slots have the required free links. All nodes sends a packet back to the requesting node to notify which slots, if any, that have been allocated. When the requesting node has received the answers, it decides if it is satisfied with the number of allocated slots. If not, it sends a release packet. Otherwise, it can start using the reserved slots immediately. If so, the node notifies other nodes by broadcasts the details of the RTVC. It should also send a release packet if more slots than needed were allocated. A node wishing to set up an RTVC can thus "borrow" slots from other nodes. A further, more detailed, description of this is found in [6].

5 Implementation aspects

T_{tcma} is the time required to complete network arbitration according the TCMA protocol. This also sets the minimum possible slot length for the network. T_{tcma}, scales for increasing network length and number of nodes as follows:

$$T_{tcma} = T_{collection} + T_p + T_{selection} + T_{distribution} \qquad (1)$$

and is explained below. The master requires a non-infinite time for processing the collected requests to select which requests may be sent. Part of this processing is sorting the incoming requests. This can be done as they arrive by checking them bit by bit and thus the sorting time is incorporated in the time for the collection phase, $T_{collection}$. When requests have been sorted, the list of requests has to be traversed to select which may be sent. The time for this is denoted as $T_{selection}$ and is assumed to be $N \cdot 30$ ns. Propagation delay, T_p, is part of arbitration since the master depends on feedback, i.e., the requests from the other nodes in the collection phase. *L* is the total length of the ring. *P* is the propagation delay through the optical fibre and is assumed to be 5 ns/m. The delay through each node (approximately 1 bit time per node) is neglected in the propagation delay, thus the total propagation delay around the ring is:

$$T_p = L \cdot P \qquad (2)$$

From Figure 4 one can see the size of the fields in the control packets which lead to:

$$T_{collection} = \frac{1 + (2N + 4)(N - 1)}{C} \qquad (3)$$

and

$$T_{distribution} = \frac{10N}{C} \qquad (4)$$

where *C* is the bit rate of the links which is 800Mb/s for the Motorola Optobus and *N* is the number of nodes. The minimum slot length is plotted in Figure 6. As we can see from the figure, the minimum slot length increases with increased network length and number of nodes. This is just as can be expected from a network that requires feedback in its medium access protocol.

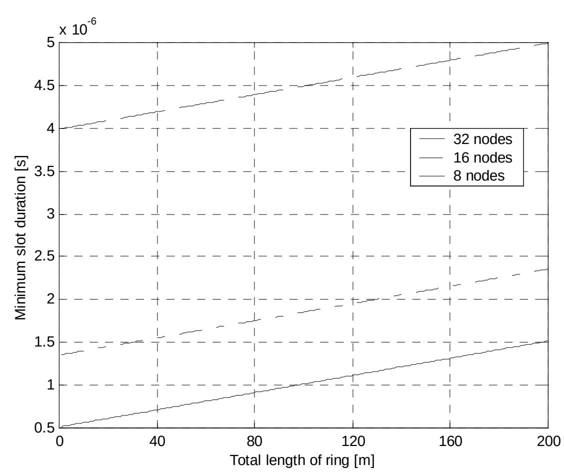

Figure 6: The diagram shows the relation between total network length and minimum slot length, for three different numbers of nodes.

5.1 Worst-case analysis

From the point in time when a packet is generated the delay until the packet may be sent consists of two components. The first component is the delay because other packets in the network are more urgent or have equal priority but are destined further and thus have priority. The second component is when no other packets that are equally or more urgent on other nodes, i.e., the packet is next in line to be sent. Only the second part of the delay is analysed below. This component of the delay is referred to as access to the network latency.

Two cases for worst case access to the network latency are presented: when the packet is destined M=1 hops and when the packet is destined M=N-1 hops. These two cases are explained in Figure 7. As we can see in the figure both latencies are minimised by sending the control phase (CP) packets as close as possible to the next slot. The vertical lines denote the end and beginning of slots. Since the clock is interrupted at the end of each slot, it is impractical to have the CP packets be sent at the end of one slot and continue into the next slot. As can be seen in the figure, slot length also affects the latency.

When a node is master, it is always granted transmission as long as no other node has a packet with a lower priority level (closer to deadline). It is assumed that a node is master every N slots and that no other node has more urgent packets. The equations for the latency is presented below:

$$T_{latency} = M \cdot T_{slot} + T_{cpp\_skew} \tag{5}$$

where

$$T_{cpp\_skew} = T_{tcma} + T_n \tag{6}$$

and

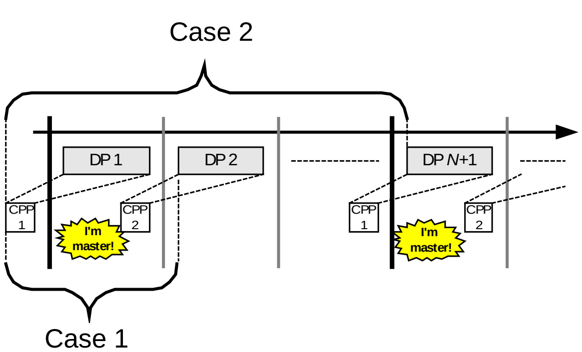

Figure 7: Worst case latency for the two cases. 1: *M*=1 hops and 2: *M*=*N*-1 hops. Control phase packets are denoted CPP and imply both control and distribution phases. Data packets are denoted DP. Note that the diagram is schematic and not to scale.

$$T_n = \frac{N-1}{B} \tag{7}$$

T_{slot} is the slot length, and T_n is the total extra delay incurred to the control packet, compared to the data packet, when passing through N-1 nodes. The extra delay is assumed to be one bit-time in each node. T_{cpp_skew} is the time from the start of the control phase until the start of the data packet. This can be seen in Figure 7. Remember that the control phase for the current slot occurred in the previous slot.

As can be seen in Figure 7 the latency depends on how far the node wishes to transmit, denoted M. The importance of keeping note of M is that a node may not transmit to another node that is "past" the master. This is because the clock signal is interrupted at the master. Transmissions destined for downstream neighbours (M=1) are always possible whereas the chance to be able to send decreases with increased number of hops. This explains the M-term in Equation 5.

Equation 5 is plotted in Figure 8. T_{slot} is chosen to be 5 µs, see the previous section on minimum slot length. As can be seen in the figure, there is a large difference in latency depending on the destination of the transmission. A user of the network can therefore gain in performance by carefully optimising algorithms to take advantage of this property, which is, sending to the next neighbour. The reason that the plots in the figure are almost horizontal, especially the plot for M=N-1 is that the slot delay ($M \cdot T_{slot}$) is a very small part compared to the rest of the arbitration delay (T_{cpp_skew}).

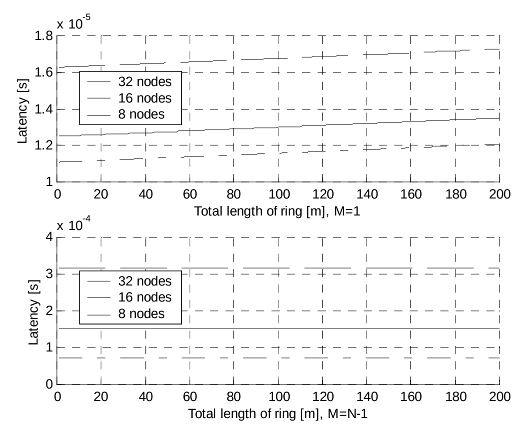

Figure 8: The worst case latency for two cases: *M*=1 hops and *M*=*N*-1 hops. The plots are made using a slot length (T_{slot})of 5 µs, which seems to be a realistic choice, see also Figure 6.

5.2 Simulation analysis

In addition to worst case performance analysis reported on in the previous section, an average case performance analysis for best effort traffic is also presented. The analysis is done by discrete time computer simulation. Networks of 8, 16, 32 and 64 nodes were simulated.

Packet in the system have soft real time constraints, thus are not given any guarantee, at generation, that it will be sent in a timely fashion. The packet is given a relative deadline at generation which is decremented each time slot. The packet is then queued until it is either sent successfully or, when deadline reaches zero, is deemed lost and removed from the queue. We also call this type of communication, best effort messages [10]. The user sending packets with a best effort service should not require any guarantees.

Some further assumptions for the simulations:

- Messages are one packet long and take one time slot to send. The term packet and message is therefore used synonymously.
- Uniform traffic is assumed, i.e., all nodes have equal probability of message generation and uniformly distributed destination addresses. This implies that, on average, it is theoretically possible to transmit two packets each slot, since the packet on average is destined "half way" around the ring, which is $N / 2$ hops. However, this disregards protocol effects which lowers the average utilisation which we will see later. An example of protocol effect is that a node may not send past the master since the clock is interrupted there. The pipelined ring topology of the network suggests that it be very effectively utilised when traffic is mostly destined one hop to the next neighbour such as in some types of radar signal processing [11], [12]. If this special mode of traffic were simulated, which it is not, the effect would simply be higher throughput because of aggregation, i.e., several packets would be sent during one slot.
- Messages were generated according to a poisson process and all messages were of single destination type.
- The deadline of all best effort packets is set at generation to 800 slot times ahead, which would equate to a deadline of 4ms with a slot time of 5 µs.
- Physical effects such as the propagation delay through optical fibre and the time required to detect the end of a slot was assumed to be less than one slot time and is therefore neglected.
- Message latency is defined as the time elapsed from the moment a message is generated until the entire message is received in the receiver.
- Infinite size of message queues is assumed.
- The simulator is run for a total of 100 000 slot times and starts to log statistics at 20 000 slot times.
- The "total packet generation intensity" is the generation intensity for the network regarded as a whole, not of the individual nodes.

Figure 9 shows the best effort packet latency for varying levels of network throughput. At a useful level of packet intensity the network has an average throughput of approximately 1.6 packets per slot. This is 60% better than the theoretical limit of one packet per slot for networks without spatial reuse. Similar results for varying number of nodes are obtained. Figure 9 shows up to which level of throughput the network may be useful. For a clearer view of throughput see also Figure 10.

Figure 10 shows the packet throughput against packet generation intensity. Throughput has a linear relation to packet intensity up to the point of saturation, which can be seen in the figure as the point on the plot where throughput starts to decrease with increasing packet intensity. As packet intensity passes the point of saturation it becomes increasingly difficult for the packet transmission scheduler to effectively utilise the bit rate of the network. This is because it is always increasingly difficult to schedule packets the further their destination

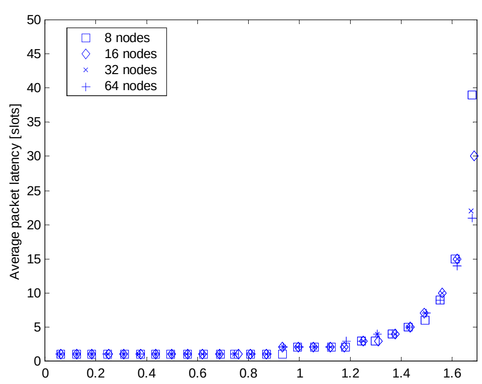

Figure 9. Best effort packet latency vs. throughput for a varying number of nodes.

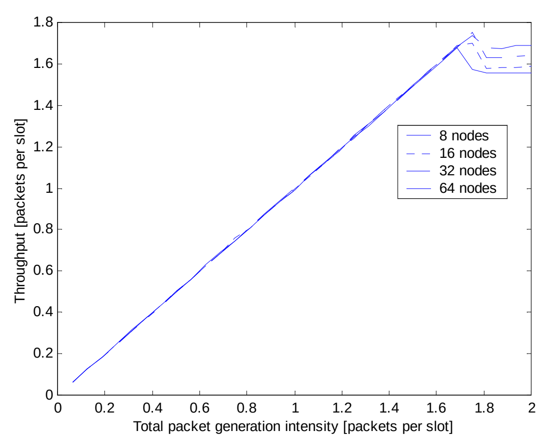

Figure 10: Total throughput of best effort packets vs. packet generation intensity.

[2]. When the network is saturated each node will tend to always contain far destined packets with short deadlines. These cannot always be scheduled together with the shorter destined packets because of conflicting links. When two packets have the same priority level the packets with the further destination has priority. Thus slot utilisation will decrease which is expressed in the decline of the throughput after the “summit” in the figure. The simulation shows that the utilisation of slots does not perform as well as may be theoretically expected (two packets transmitted per slot, because of pipelining). This is attributed to the policy of using deadline to selecting packets for transmission. The policy is not as bandwidth conserving as, e.g., the FWS policy [2] which can achieve higher average throughput. This is the cost of having good support for real-time traffic.

Figure 11 shows packet throughput and average packet latency for each destination distance plotted against packet generation intensity. The simulation is for 16 nodes. Concluded from this simulation is that TCMA treats packets fairly regardless of the distance to the destination even when the network approaches and is saturated. Observe that there are 15 (for N-1 distances to destination) plots for latency but that these overlap and appear as one.

6 Conclusions

This paper presents an optical ring network medium access protocol that globally optimises packet timing constraints. Simulation results of the medium access protocol have been presented. The protocol is shown to be fair even when the network is saturated. The network with the presented protocol is suitable for application such as in embedded systems, e.g., for use as interconnection network in a radar signal processing system, or as a high performance network for a LAN. Also worth mentioning is that the network can be built today using fibre-optic off-the-shelf components

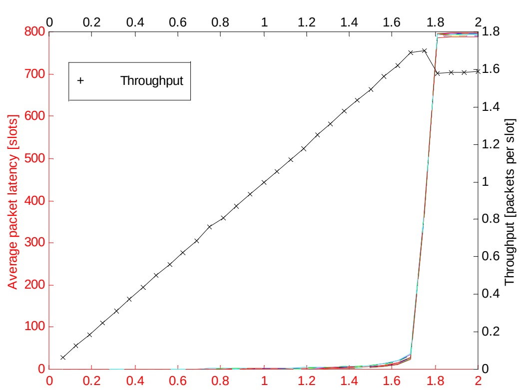

Figure 11. Packet throughput vs. packet generation intensity for 16 nodes. The other curves represent latency for the (15) different distances between source and destination.

Acknowledgement

This work is part of M-NET, a project financed by ARTES: A Real-Time network of graduate Education in Sweden.

References

[1] M. Jonsson, “Two fibre-ribbon ring networks for parallel and distributed computing systems,” *Optical Engineering*, vol. 37, no. 12, pp. 3196-3204, Dec. 1998.

[2] P. C. Wong and T.-S. P. Yum, "Design and analysis of a pipeline ring protocol," *IEEE Transactions on communications*, vol. 42, no. 2/3/4, pp. 1153-1161, Feb./Mar./Apr. 1994.

[3] M. Conti, E. Gregori, L. Lenzini, “DQDB under heavy load: performance evaluation and fairness analysis,” *INFOCOM '90, Proceedings of the 9th Annual Joint Conference of the IEEE Computer and Communication Societies*, vol.1, pp. 313 –320, 1990.

[4] B. Raghavan, Y.-G. Kim, T.-Y. Chuang, B. Madhavan, A.F.J. Levi, “A gigabyte-per-second parallel fiber optic network interface for multimedia applications,” *IEEE Network*, Volume: 13 Issue: 1 , Jan.-Feb. 1999 pp. 20 –28.

[5] N.J. Boden, D. Cohen, R.E. Felderman, A.E. Kulawik, C.L. Seitz, J.N. Seizovic, Wen-King Su, “Myrinet: a gigabit-per-second local area network,” *IEEE Micro*, Volume: 15 Issue: 1, Feb. 1995 pp. 29 –36.

[6] M. Jonsson, C. Bergenhem, and J. Olsson, "Fibre-ribbon ring network with services for parallel processing and distributed real-time systems," *Proc. ISCA 12th International Conference on Parallel and Distributed Computing Systems (PDCS-99)*, Fort Lauderdale, FL, USA, Aug. 18-20, 1999, pp. 94-101.

[7] C. Bergenhem and J. Olsson, “Protocol suite and demonstrator for a high performance real-time network,” *Master thesis, Centre for Computer Architecture (CCA), Halmstad University, Sweden*, Jan. 1999.

[8] D. Bursky, "Parallel optical links move data at 3 Gbits/s," *Electronic Design*, vol. 42, no. 24, pp. 79-82, Nov. 21, 1994.

[9] C. Bergenhem, “A demonstrator for a CC-FPR network,” *Techical report 0010, School of Information Science, Computer and Electrical Engineering (IDE), Halmstad University, Sweden,* Feb. 2000. In Swedish.

[10] K. Arvind, K. Ramamritham, and J. A. Stankovic, “A local area network architecture for communication in distributed real-time systems,” *Journal of Real-Time Systems*, vol. 3, no. 2, pp. 115-147, May 1991.

[11] M. Jonsson, B. Svensson, M. Taveniku, and A. Åhlander, "Fiber-ribbon pipeline ring network for high-performance distributed computing systems," *Proc. International Symposium on Parallel Architectures, Algorithms and Networks (I-SPAN'97)*, Taipei, Taiwan, Dec. 18-20, 1997, pp. 138-143.

[12] M. Taveniku, A. Åhlander, M. Jonsson, and B. Svensson, "The VEGA moderately parallel MIMD, moderately parallel SIMD, architecture for high performance array signal processing," *Proc. 12th International Parallel Processing Symposium & 9th Symposium on Parallel and Distributed Processing (IPPS/SPDP98)*, Orlando, FL, USA, Mar. 30 – Apr. 3, 1998, pp. 226-232.