Transactions on Information and Communications Technologies vol 3, © 1993 WIT Press, www.witpress.com, ISSN 1743-3517

Performance evaluation of distributed transputer networks using high-performance communication links

D. Popović, H. Westphal

University of Bremen, Institut of Automation Technology, FB1, NW1, Kufsteiner Str., W-2800 Bremen, Germany

Abstract. Interface Standards (IS) [9] become available for the interconnection of distributed high performance transputer systems; for instance the Fibre Distributed Data Interface (FDDI) and the High Performance Parallel Interface (HIPPI). The FDDI with 100 Mb/s backbone speed is an exellent Interface Standard for a Transputer Link Interface (TLI) to interconnect distributed 100 Mb/s Inmos IMS T9000 links [13, 12]. For a Gigabit communication channel the high performance point-to-point HIPPY is a 1.6 Gb/s Interface Standard [7, 8]. In the past the communication links had represented the performance bottlenecks of distributed transputer systems. However, now having enormous bandwidth the communication bandwidth is no longer a constraint. For the application of multitransputer systems in solving complex computational problems for distributed automatic control and in order to meet the requirements on the underlying communication system the paper gives an answer to the question - do gigabit transputer links for distributed transputer systems represent just another step in the evolutionary process of greater bandwidth systems ?

1 Introduction

Distributed Computer Control Systems (DCCS) [11] are pertinent part of modern automation technology, particularly for real-time process control. Engineering, scientific and commercial tasks basically require parallel processing, because the sequential processing is not able to meet the requirements in many modern application fields like plant automation, image processing, high-speed communication, intelligent systems etc.. Such problems can more efficiently be solved by using the transputer-based distributed computer systems, that render a high computation performance through modular, multiple architecture, and parallel multiprocessing capabilities.

Transputers use within the multinode supercomputer systems brings, due to the short distance between the directly connected transputers, effectively no data transmission problems. In DCCS, however, data transmission takes place across a serial system network so that the attachement of individual transputers to the network brings, mainly due to the asynchronous handshake link protocol, enormous problems, this being the consequence of *signal delay time.*

When the application field requires a communication bandwidth in the megabit world ($n*100Mb/s$), between the remote transputer stations, FDDI gives the technology for the remote tranputer link (**section 2.1**). We are moving into an era of gigabit per second speeds for I/O Interfaces and networks. For application fields that require a communication bandwidth in the gigabit world ($n*1.6Gb/s$) HIPPI gives the technology for remote transputer links (**section 2.2**).

Evaluation of latency and bandwidth shows that one way to hide latency is to use parallelism such that while one process is waiting for a response, another process, which does not depend upon this response, may proceed with its processing (**section 2.3**). This procedure can be realized by using the virtual channel technology of the IMS T9000 and the FDDI or the HIPPI protocol (**section 3**) for a distributed transputer link. But it is clear that more bandwidth is of no use at all in speeding up the transmission of a message, it is the latency of the link that dominates the time to deliver the message (**section 4.1**). That is we can benefit from more bandwidth. The quantity of parallel processes, that can be implemented using the T9000 and the remote transputer link, and a Transputer Process Response Time (TPRT) evaluation is given in (**section 4.2**). A critical Bandwidth C_{crit} is defined for the boundary of a latency-, or bandwidth-limited transputer link (**section 4.3**).

2 System Elements

In this section the fundamental system elements (FDDI, HIPPI and T9000) will briefly be described.

2.1 FDDI

In distributed Transputer Systems with application fields that require a communication bandwidth in the megabit world ($n*100Mb/s$), [13] proposes a TLI, based on the Inmos IMS T9000 [14] and the FDDI [6, 4]. Using FDDI as the technology for remote transputer stations one serial link of a T9000 in a transputer station is interconnected to a serial link of a T9000 in a remote transputer station. **Fig. 1** shows transputer system T_1 and the remote transputer system T_2 exchanging data by means of transputer link interface TLI_1, a Fibre Optic Link FOL and transputer link interface TLI_2. TLI_1 and TLI_2 act as a remote communication link and the distance between two TLIs is up to 40 km using single mode fibre [5, 1]. The FDDI protocol [16] which take advantage of the TLI is processing up to the second layer, the Media Access Control (MAC) [4] layer, of the seven layer model, extracting a minimun amount of network layer

Transactions on Information and Communications Technologies vol 3, © 1993 WIT Press, www.witpress.com, ISSN 1743-3517

information. The performance evaluation in [13] shows the bandwidth and the response time characteristics (from $13\mu s$ up to $1ms$) of the FDDI TLI.

2.2 HIPPI

For distributed Transputer Systems with application fields that require a communication bandwidth in the gigabit world ($n * 1.6Gb/s$), HIPPI [7, 8] gives a technology. HIPPI-PH is a high-performance point-to-point interface standard for transmitting digital data in parallel between data processing equipment. The maximum distance supported is $25m$ with copper twisted-pair cables. Using HIPPI for remote transputer links with a distance beyond $25m$ leads to the Serial-HIPPI-Specification [2], an extender for transmitting digital data serially between HIPPI-PH nodes. Multiple serial links of T9000s in a transputer station are interconnected to multiple serial links of T9000s in a remote transputer station. The distance between two transputer stations, interconnected with the serial HIPPI, is up to $10km$ using single mode fibre and the Serial-HIPPI baud rate is $1.2Gb/s$ [2]. The $1.6Gb/s$ HIPPI is supported by two Serial-HIPPIs in parallel. Serial-HIPPI defines the parallel interface to HIPPI-PH, the serial electrical interface appropriate for short links on coaxiale cable, and the serial optical interface appropriate for long links on optical fiber as shown in **Fig. 1**.

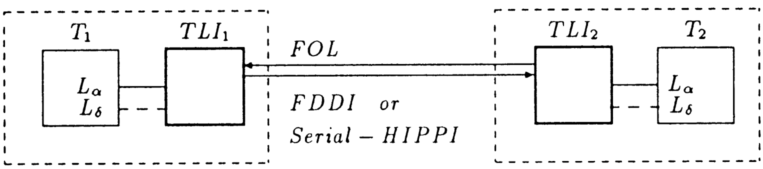

FOL : Fibre Optic Link
T : Transputer System
TLI : Transputer Link Interface
FDDI : Fibre Distributed Data Interface
HIPPI : High Performance Parallel Interface

Fig. 1. A typcal TLI application

2.3 T9000

IMS T9000 [14] is provided by four $100MBaud$ serial links L_α–L_δ, each serial link having a packet based link protocol, supporting a total bidirectional data communication bandwidth of $80MBaud$. The links are asynchronous, the receiving device synchronices the incomming data and devices with different processor speeds can communicate. Processes can intercommunicate via the special channels as shown in **Fig. 2**. For instance, the process Pr_A and Pr_n can intercommunicate via the Virtual Channel Processor (VCP) that includes the interfaces C_A, C_B and C_n, known as *'virtual channels'*, that share the link L_α between individual IMS T9000 transputers. In this way, any number of processes can

Transactions on Information and Communications Technologies vol 3, © 1993 WIT Press, www.witpress.com, ISSN 1743-3517

communicate with any number of transputers by using the VCPs. Each process runs until it has completed its action or its scheduled time is over.

Messages, sent across a link, are split up into packets (s. **Fig. 3**) that are interleaved during the transmission. Every packet requires a header for identification of its destination process. The transputer sends the first packet of a message from a process and waits, before sending the next packet of the message, for an *acknowledgement* confirmation (ACK) from the receiving processor. This repeats until the last packet of the message has been sent and acknowledged. Virtual channels are created in pairs to form a 'virtual link' in which the acknowledgements are simply sent back along the second channel of the virtual link. Messages and acknowledgements from other virtual links can be sent while waiting for an acknowledgement on a virtual link. This ensures that a single virtual link cannot monopolize a physical link.

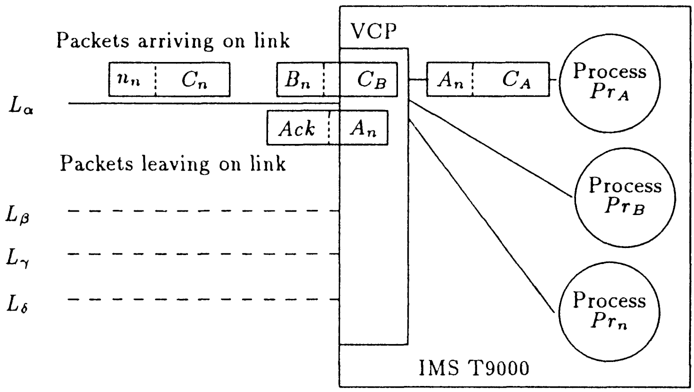

Fig. 2. parallelism to hide latency

3 Hierarchical Protocol Architecture of Communication Links

The communications links can basically be described at different hierarchical level protocols. In **Fig. 3** examples are given of transputer *high level* and FDDI and HIPPI *low level* protocol.

The highest level of protocol, the *message level* (M), consists of data or messages that the user sends, from one *process* to another, through a physical or virtual transputer channel (s. **Fig. 2**).

At the *packet level* of the protocols the packets, belonging to a number of different channels, can be sent interleaved via the same transputer link. Every packet has a header, defining the destination address, followed by the data bytes (0 to 32 data bytes) and, closed by an *end of packet* or *end of message*

Transactions on Information and Communications Technologies vol 3, © 1993 WIT Press, www.witpress.com, ISSN 1743-3517

token. This simple protocol supports messages of any length and also enables the synchronization at the message level.

When using the FDDI frame [4], the Preamble (PA) is followed by the Starting Delimiter (SD), the Frame Control (FC), the Destination Address (DA), e.g., $Adr.TLI_2$, the Source Address (SA), e.g., $Adr.TLI_1$, the Information (INFO, 0 to about $4KByte$, e.g., **transputer packets**), the Frame Check Sequence (FCS), the End Delimiter (ED), and the Frame Status (FS).

When using the HIPPI frame [8], the $64bit$ header is followed by a data area (0 to about $4GByte$ in one HIPPI-FP packet, e.g., **transputer packets**). A $64bit$ word is transferred on every $40ns$ cycle; words are grouped into bursts and bursts are grouped into packets. Bursts contain $256words$ (or fewer), and packets consist of 1 or more bursts.

In accordance to the use of the TLI as a remote transputer link, we do not have to discuss the operation of the higher level protocols of the seven layer model, that allow services provided and required by MAC, to operate correctly (e.g., MAC to LLC [3]).

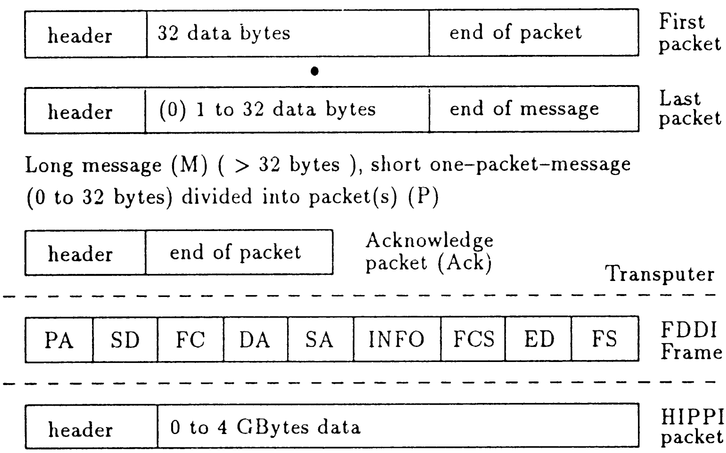

Fig. 3. high level transputer-, low level FDDI-, low level HIPPI protocol

4 Systems Parameter Evaluation

4.1 Latency and Bandwidth evaluation due to the speed of light

In DCCS data transmission takes place across a serial system network with a signal delay time, this being the consequence of the limitation of the speed of light. Let us examine the effect of FDDIs and HIPPIs latency and bandwidth due to the speed of light. There are a few key parameters of interest in any data network system [10]. These are:

C = Capacity of the network (Mb/s)
b = Number of bits in a data packet
L = Length of the network (km)
We think of the network as a remote transputer link. Combining these three parameters gives a single critical system parameter a, which is defined as:

$$a = \frac{5LC}{b}. \tag{1}$$

The factor 5 appearing in the equation is simply the approximate number of microseconds it takes light to move 1000 m ($5,085\mu s/km$). This parameter is the ratio of the latency of the link (i.e. the time it takes energy to move from one end of the link to the other) to the time it takes to pump one packet into the link. The parameter a measures how many packets can be pumped into one end of the link before the first bit appears at the other end.

For the calculation of the **FDDI Latency and Bandwidth** we assume a stable operation state in which each TLI subsequently receives a token, transmitts a frame, and a token. A one–packet–message, composed of a transputer data and Ack packet ($312bit$), as INFO in the FDDI frame ($224bit$, without INFO), gives a FDDI frame length of $536bit$, i.e., $5,36\mu s$ transmission time, a ten packet message of $3344bit$ and, a hundred packet message of $31424bit$ (**s. section 3**). The Token time ($0,88\mu s$) consists of the time required to transmit a token ($24bit$) and its preamble ($64bit$). Token capture and retransmission time is assumed to be on the order of a few bits and not significant to MAC timer calculation [4]. For the one, ten and hundred packet transputer message encapsulated in the FDDI protocol the calculation of the ratio a is shown in **table 1**. **Table 2** shows the ratio $1/a$ (and its diagram) that gives the "FDDI length" of the transputer message to the length of the link.

For the calculation of the **HIPPI Latency and Bandwidth** we assume the same stable operation state as we used for the FDDI. A one–packet–transputer–message ($312bit$), as data in the HIPPI packet ($64bit$, without data), gives a HIPPI frame length of $376bit$, i.e., $0.235\mu s$ transmission time, a ten packet message of $3184bit$ and, a hundred packet message (8 bursts) of $31712bit$ (**s. section 3**). For the one, ten and hundred packet transputer message encapsulated in the HIPPI protocol the calculation of the ratio a is shown in **table 3**. **Table 4** shows the ratio $1/a$ (and its diagram) that gives the "HIPPI length" of the transputer message to the length of the link.

Note that in our one packet transputer message one, in our ten packet transputer message ten, and in our hundred packet transputer message hundred logically parallel transputer processes are encapsulated in the FDDI or the HIPPI protocol. Note the enormous range for the parameter a. At one extreme, namely FDDI, using the one packet message and a distance of 1000 m, it is, e.g., approximately 1 (**s. Table 1**), while at the other extreme, namely HIPPI, it is approximately 21 (**s. Table 3**) using the same configuration. We see that the parameter a grows dramatically when we introduce gigabit links. So we must ask ourselves if remote transputer links made out of gigabit links are different in some fundamental way from those made out of megabit links.

Transactions on Information and Communications Technologies vol 3, © 1993 WIT Press, www.witpress.com, ISSN 1743-3517

propagation delay/Packet Tx time				link buffering/TLI buffer					
C	b	t	a	bl	m	$1/a$	TLI_1	m	TLI_2
100	536	0.005	$9.33 * 10^{-4}$	2	1	$1.07 * 10^{3}$			
		0.050	$9.33 * 10^{-3}$		10	$1.07 * 10^{2}$			
		0.500	$9.33 * 10^{-2}$		100	$1.07 * 10^{1}$			
		5.000	$9.33 * 10^{-1}$		1000	$1.07 * 10^{0}$			
		50.000	$9.33 * 10^{0}$		10000	$1.07 * 10^{-1}$			
		200.000	$37.31 * 10^{0}$		40000	$2.68 * 10^{-2}$			
	3344	0.005	$1.50 * 10^{-4}$		1	$6.7 * 10^{3}$			
		0.050	$1.50 * 10^{-3}$		10	$6.7 * 10^{2}$			
		0.500	$1.50 * 10^{-2}$		100	$6.7 * 10^{1}$			
		5.000	$1.50 * 10^{-1}$		1000	$6.7 * 10^{0}$			
		50.000	$1.50 * 10^{0}$		10000	$6.7 * 10^{-1}$			
		200.000	$5.98 * 10^{0}$		40000	$1.7 * 10^{-1}$			
	31424	0.005	$1.60 * 10^{-5}$		1	$6.28 * 10^{4}$			
		0.050	$1.60 * 10^{-4}$		10	$6.28 * 10^{3}$			
		0.500	$1.60 * 10^{-3}$		100	$6.28 * 10^{2}$			
		5.000	$1.60 * 10^{-2}$		1000	$6.28 * 10^{1}$			
		50.000	$1.60 * 10^{-1}$		10000	$6.28 * 10^{0}$			
		200.000	$6.36 * 10^{-1}$		40000	$1.57 * 10^{0}$			

capacity (C) Mb/s
packet length (b) bits
propagation delay (t) microsec
ratio (a)

FDDI bit length (bl) meter
link distance (m) meter
ratio ($1/a$)
transputer link interface (TLI)

Table 1. propagation delay / one,- ten,- hundred transputer packet FDDI-Tx time
Table 2. link buffering of the FDDI frame / TLI buffering

To see the effect of this, let us consider the following scenario. Assume TLI_1 (**Fig. 1**) wish to send a one packet transputer message (536 bit, using FDDI; 376 bit, using HIPPI) to TLI_2.

If we use the $100Mb/s$ FDDI and a distance between the transputer stations of 1000 m, then, as shown in **Table 2**, the first bit of this transmission will arrive at TLI_2 after approximately all bits (all exept approximately 7% of the message) have been pumped into the FOL. Thus we see that the FOL is buffering roughly all of the message.

If we use the $1.6Gb/s$ HIPPI and the same configuration then, as shown in **Table 4**, we see the entire one packet message as a small pulse moving down the link. The pulse occupies roughly only 0.047 of the FOL "buffer".

Clearly, if we have a higher-speed link, the time to transmit our one packet message could be reduced. It is now clear that more bandwidth is of no use at all in speeding up the transmission of the message; it is the latency of the link that dominates the time to deliver the message. That is, we can benefit from more bandwidth.

propagation delay/Packet Tx time				link buffering/TLI buffer					
C	b	t	a	bl	m	$1/a$	TLI_1	m	TLI_2
1.6	376	0.005	$2.13 * 10^{-2}$	2	1	$4.7 * 10^{1}$			
		0.050	$2.13 * 10^{-1}$		10	$4.7 * 10^{0}$			
		0.500	$2.13 * 10^{0}$		100	$4.7 * 10^{-1}$			
		5.000	$2.13 * 10^{1}$		1000	$4.7 * 10^{-2}$			
		50.000	$2.13 * 10^{2}$		10000	$4.7 * 10^{-3}$			
	3184	0.005	$2.51 * 10^{-3}$		1	$4 * 10^{2}$			
		0.050	$2.51 * 10^{-2}$		10	$4 * 10^{1}$			
		0.500	$2.51 * 10^{-1}$		100	$4 * 10^{0}$			
		5.000	$2.51 * 10^{0}$		1000	$4 * 10^{-1}$			
		50.000	$2.51 * 10^{1}$		10000	$4 * 10^{-2}$			
	31712	0.005	$2.52 * 10^{-4}$		1	$3.96 * 10^{3}$			
		0.050	$2.52 * 10^{-3}$		10	$3.96 * 10^{2}$			
		0.500	$2.52 * 10^{-2}$		100	$3.96 * 10^{1}$			
		5.000	$2.52 * 10^{-1}$		1000	$3.96 * 10^{0}$			
		50.000	$2.52 * 10^{0}$		10000	$3.96 * 10^{-1}$			

capacity (C) Gb/s
packet length (b) bits
propagation delay (t) microsec
ratio (a)

HIPPI bit length (bl) meter
link distance (m) meter
ratio ($1/a$)
transputer link interface (TLI)

Table 3. propagation delay / one,- ten,- hundred transputer packet HIPPI-Tx time
Table 4. link buffering of the HIPPI packet / TLI buffering

4.2 Response time and bandwidth evaluation versus load

The *real–time response time and bandwidth evaluation* is based on estimation of the time delay of the system and on the throughput analysis for transmitting of transputer packets over the FDDI [15] and the HIPPI. The delay (from the time a packet P from a message M of a process Pr is generated at one transputer, until it reaches the remote transputer process, and until the acknowledge (Ack) of this packet reaches the origin process) is the sum of individual time–delays encounted at each transmission step; we call it the Transputer Process Response Time (T_{TPRT}). Each of the step mentioned can be modeled, along with its distribution function, using queuing theory [15]. In this way, the process–to–process response time and bandwidth of the remote transputer link can be estimated.

Let us now consider the response time for transputer one,- ten,- and hundred packet messages encapsulated in the FDDI or the HIPPI frame. If, as usual, we let ρ denote the system utilization factor, then

$$\rho = \lambda \frac{b}{C}, \tag{2}$$

where λ is the arrival rate (messages per microsecond) and b the number of bits in a packet and C the channel capacity (Mb/s). Because of the realtime ability of

Transactions on Information and Communications Technologies vol 3, © 1993 WIT Press, www.witpress.com, ISSN 1743-3517

the remote transputer link we assume the boundary condition $0 \leq \rho < 1$. **Table 5** shows for this assumption ($\rho < 1$) the maximum arrival rate λ_{max} of transputer messages encapsulated in the FDDI or the HIPPI frame and the approximate maximum number of Virtual Parallel Transputer Messages ($VPTM_{max}$) and Virtual Parallel Transputer Packets ($VPTP_{max}$) (s. **section 2.3**) that can be implemented using the T9000.

λ_{max}, $VPTM_{max}$ and $VPTP_{max}$					
	C	b	λ_{max}	$VPTM_{max}$	$VPTP_{max}$
FDDI	100	536	$1.86 * 10^{-1}$	186567	186567
		3344	$2.99 * 10^{-2}$	29904	299040
		31424	$3.18 * 10^{-3}$	3182	318200
HIPPI	1600	376	$4.25 * 10^{0}$	4255319	4255319
		3184	$5.02 * 10^{-1}$	502512	5025120
		31712	$5.04 * 10^{-2}$	50454	5045400

arrival rate λ_{max} (messages per microsecond)
$VPTM_{max}$ (messages per second)
$VPTP_{max}$ (packets per second)

T_{TPRT}					
d	T_{TPRT}	d	T_{TPRT}	d	T_{TPRT}
1	16	5	56	10	106
	44		84		134
	321		365		415
	11		51		101
	12		52		102
	30		70		120

distance d (km)
T_{TPRT} (microseconds)

Table 5. λ_{max}, $VPTM_{max}$ and $VPTP_{max}$
Table 6. T_{TPRT}

The transputer process response time T_{TPRT} (microseconds) of the remote transputer link (i.e. the time from when the message is send by a TLI until the last bit of the message appears at the output of the link, and the acknowledge arrives at the sending TLI), is given by

$$T_{TPRT} = \frac{b}{C(1-\rho)} + \tau, \qquad (3)$$

where τ is the propagation delay (i.e. the channel latency) in microseconds.

Let us ask ourselves if gigabit links actually help in reducing the response time, T_{TPRT}. **Table 6** shows the T_{TPRT} (in microseconds) for different transputer station distances (d, in km) and for the one,- ten,- and hundred packet transputer message encapsulated in the FDDI ($100 Mb/s$) or HIPPI ($1.6 Gb/s$) protocol. For a one packet FDDI transputer message ($536 bit$, s. **Table 5**) and a transputer system distance of $1 km$ (i.e. a FOL distance of $2 km$, $10 \mu s$ propagation delay) the T_{TPRT} is $16 \mu s$.

Fig. 4 (**Fig. 5**) (**Fig. 6**) shows the T_{TPRT} for the one,- ten,- and hundred packet transputer message versus the system load ρ for the FDDI and the HIPPI remote transputer link and a transputer system distance of $1 km$ ($5 km$) ($10 km$). For the initial values refer to **Table 6** with the distance (d) of 1, 5 and $10 km$. The extremely high load situation is to be avoided for realtime ability. We note that when we go from $100 Mb/s$ to $1.6 Gb/s$, we see for a transputer system distance of approximately less than $5km$ a improvement. As far as response time is concerned, gigabits do help here.

4.3 Critical bandwidth C_{crit}, a boundary for a latency or bandwidth limited link

There are two components, as can be seen from (**Eq. 3**), making up the response time, namely, the queuing-plus-transmission time delay (the first term in the equation) and the propagation delay (τ). In accordance with the boundary condition $0 \leq \rho < 1$ (realtime ability) we do not have to discuss the queuing delay. We discuss the relative size of each of these and we reffer to regions of bandwidth-limited and latency-limited links (**s. Table 2,4**). Let us now make this concept more precise. We choose to define a sharp boundary between these two regions. In particular, we define this boundary to be a place where the two terms in our equation are exactly equal, namely, where the propagation delay equals the transmission delay. From **Eq. 3** we see that this occurs when the bandwidth of the link takes on the following critical value,

$$C_{crit} = \frac{b}{(1-\rho)\tau}. \tag{4}$$

Fig. 7 (8) (9) shows this critical value of bandwidth C_{crit} – for a transputer system distance of $1km$, $\tau = 10\mu s$ ($5km$, $\tau = 50\mu s$) ($10km$, $\tau = 100\mu s$), and the maximum length (b) of the one,- (536 bit) ten,- (3344 bit) and hundred (31712 bit) packet transputer message – versus the system load ρ. Above this boundary, the system is latency limited which means that more bandwidth will have negligible effect in reducing the response time, T_{TPRT}. Below this boundary, the system is bandwidth limited which means that it can take of advantage of more bandwith to reduce T_{TPRT}. For the one packet transputer message and a transputer system distance of $5km$ the system is latency limited over most of the load range when a the FDDI ($100Mb/s$) is used; this means that for these parameters a gigabit channel is overkill so far as reducing delay is concerned. For the ten packet message and the same distance ($5km$) the system is bandwidth limited over most of the load range using the FDDI ($100Mb/s$); this means that for these parameters a gigabit channel begins to make sense so far as reducing delay is concerned. In accordance with our configuration we note further that gigabit channels make sense for relative big message sizes, but are not helpful for small message sizes. For other configurations, the critical bandwidth which defines the boundary is given from **Eq. 4**.

5 Conclusion

In this paper we have dealt with the latency and bandwidth in distributed transputer networks. One conclusion of this paper is to recognize that FDDI or HIPPI remote transputer links have forced us to deal with the propagation delay due to the finite speed of light. As we saw earlier, the propagation delay of our $2kmFOL$

is roughly 40 times (2 times) smaller than the time required to transmit a one packet transputer message into a HIPPI (FDDI) *FOL*. We must rethink a number of issues. For example, the user must pay attention to his file sizes and how latency will affect his applications. At the application level, it is important to find ways to hide the latency of the link, in order to get full advantage of the high-speed links and of the high performance processors attached to the links. One form of hide latency is to use parallelism (or pipelining) such that while one process is waiting for a response, another process, which does not depend upon this response, may proceed with its processing. This procedure, as we saw earlier, can be realized by using the virtual channel technology of the IMS T9000.

References

1. ANSI: X3T9.5. *FDDI Single-mode Fibre Physical Layer Medium Dependent (SMF-PMD)*, 1991.
2. Hewlett-Packard Labs. *Serial-HIPPI-Specification*, 1991. Serial HIPPI Implementors Group, 1991.
3. ISO 8802.2. *Logical Link Control (LLC)*, 1989. Information Processing Systems - Local Area Networks.
4. ISO 9314-2. *Information processing systems - Fibre Distributed Data Interface (FDDI) - Part 2: Token Ring Media Access Control (MAC)*, 1989.
5. ISO 9314-3. *Information processing systems - Fibre Distributed Data Interface (FDDI) - Part 3: Token Ring Physical Layer, Medium Dependent (PMD)*, 1990.
6. ISO 9341-1. *Information processing systems - Fibre Distributed Data Interface (FDDI) - Part 1: Token Ring Physical Layer Protocol (PHY)*, 1989.
7. ISO/IEC 11518-1. *Information technology, Mechanical, Electrical and Signaling Protocol Specification (HIPPI-PH)*, 1993.
8. ISO/IEC 11518-2, ANSC X3T9.3, Draft 4.3. *Framing Protocol (HIPPI-FP)*, 1993.
9. ISO/IEC JTC 1/SC 25/WG 4. *Interconnection of computer systems and attached equipment*, 1992. An Introduction to X3T9 and I/O Interface Standards.
10. Leonard Kleinrock. The latency/bandwidth tradeoff in gigabit networks. *IEEE Communications Magazine*, April 1992.
11. D. Popović and V.P. Bhatkar. *Distributed Computer Control for Industrial Automation.* Dekker, New York and Basel, 1990.
12. D. Popović and H. Westphal. Performance evaluation of an fddi-based distributed transputer network. Seminar on Transputers and Control for Indian Experts, Bordeaux, France, 1992.
13. D. Popović and H. Westphal. *Leistungsanalyse dezentraler Transputer–Systeme basierend auf FDDI Links. In: Informatik aktuell; Parallele Datenverarbeitung mit dem Transputer, R. Grebe, M. Baumann, RWTH Aachen (Hrsg.).* Springer, 1993.
14. SGS-THOMSON. *IMS T9000 Transputer*, 1992.
15. H. Westphal. Performance analysis of a distributed transputer network. University of Bremen, Institute of Automation Technology, 1992.
16. H. Westphal. Implementation of an FDDI ring. Forschungsbericht, Universität Bremen, Institut für Automatisierungstechnik, 1993.

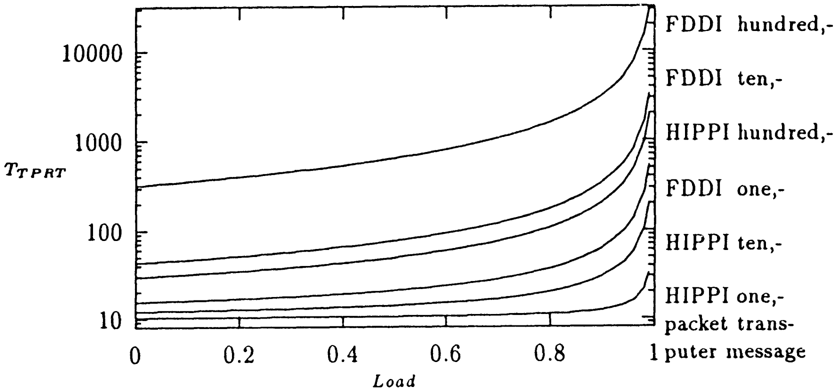

Fig. 4. T_{TPRT} (μs) versus load for a transputer system distance (d) of $1km$

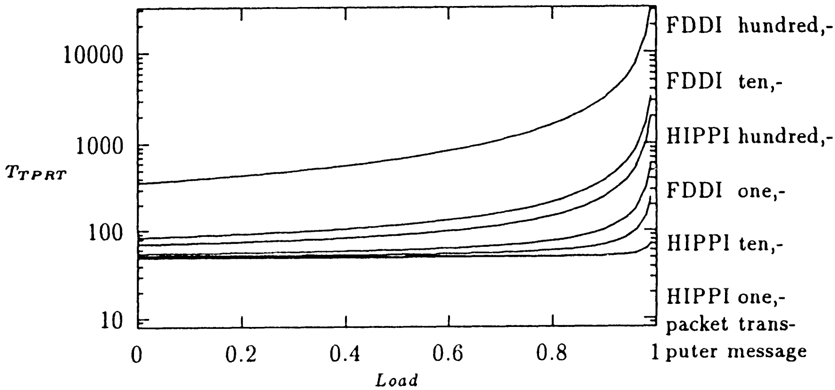

Fig. 5. T_{TPRT} (μs) versus load for a transputer system distance (d) of $5km$

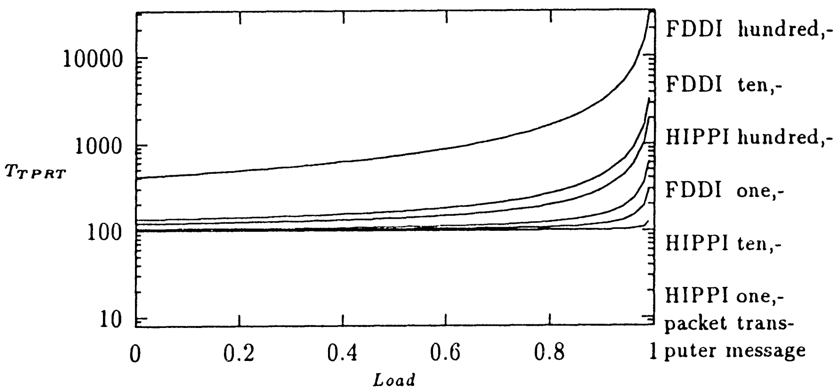

Fig. 6. T_{TPRT} (μs) versus load for a transputer system distance (d) of $10km$

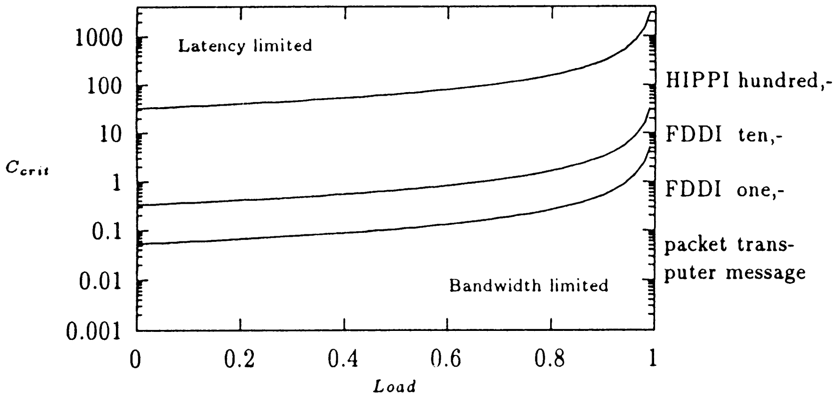

Fig. 7. C_{crit} (Gb/s) versus load for a transputer system distance (d) of $1km$

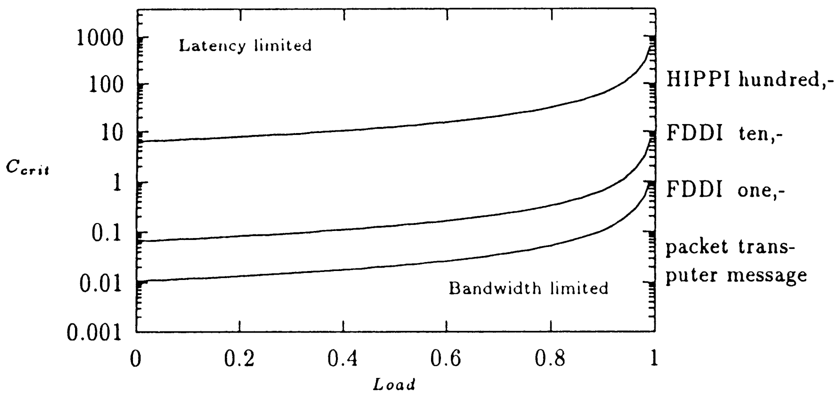

Fig. 8. C_{crit} (Gb/s) versus load for a transputer system distance (d) of $5km$

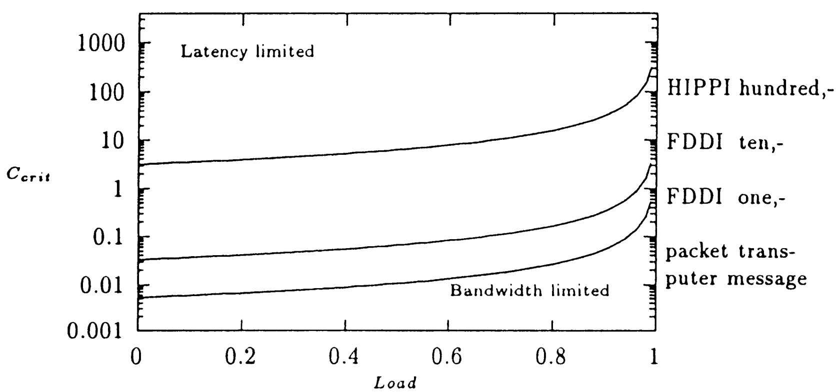

Fig. 9. C_{crit} (Gb/s) versus load for a transputer system distance (d) of $10km$