

# The Token Grid Network

Terence D. Todd, *Member, IEEE*

**Abstract**—An overwhelming majority of Local and Metropolitan Area Network products (LAN's and MAN's) are based upon linear topologies such as buses and rings [1], [2], [4], [9]. Such networks are economical for high speed operation since the station interfaces are simple and require very little transit buffering. However because of their linear structure, the total throughput is restricted by the transmission rate of the media access channels.

In this paper, a token grid network is introduced where media access is performed over a two-dimensional mesh. In the resulting system, each station is two-connected and has the same transmission hardware and small station latency as in a dual token ring [3], [4]. In the token grid however, the total system throughput may be many factors larger than that which is possible in a dual token ring. In a large  $\sqrt{N} \times \sqrt{N}$  network, the uniform load capacity is approximately  $\sqrt{N}/2$  times that of an  $N$  station dual token ring. In addition, the token grid can take advantage of communities-of-interest amongst the stations. It is possible to implement the system in such a way as to achieve robust operation in the presence of station and link failures.

## I. INTRODUCTION

MOST existing LAN/MAN products and standards are based upon linear shared media designs such as buses (e.g., IEEE 802.3 [1]) and rings (e.g., IEEE 802.5 and FDDI [2], [4]). These topologies are popular because it is possible to build a very simple and economical station attachment unit. In Ethernet for example, stations passively tap a broadcast bus medium and thus no high speed transit buffering is needed at the station. Similarly in a ring, the station interface need only buffer a small number of bits clocked at the network rate. These simple designs, coupled with the ease of controlling a linear network has led to the proliferation of many inexpensive products.

Due to the linear topology of rings and buses, the maximum throughput is severely restricted by the data rate of the shared channel. In a conventional token ring such as IEEE 802.5 [2], for example, the total throughput of the network is upper bounded by the channel data rate regardless of the number of stations. In principle, this problem may be addressed using a network topology which permits multiple concurrent packet transmissions. However, in many proposed designs the station simplicity of a bus or ring network must be sacrificed to achieve the improved performance properties. For example, in multihop or bridged networks [5], [8], high speed transit buffering and routing must be provided at each station.

Manuscript received June 18, 1993; revised March 11, 1994; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor R. Jain. This work was supported by the Telecommunications Research Institute of Ontario (TRIO).

The author was with the Distributed Systems Research Department, AT&T Bell Laboratories, Murray Hill, NJ. He is now with the Communications Research Laboratory, the Department of Electrical and Computer Engineering, McMaster University, Hamilton, Ont. L8S 4K1, Canada.

IEEE Log Number 9403693.

Although deflection routing [6] may be used to limit the buffering component, resequencing buffers are then required to recover from packet misordering due to deflections. In both cases the required station resources may be much higher than that in typical ring and bus networks.

In this paper, a network referred to as the token grid is introduced. The token grid is a multidimensional extension of the token ring [2], [4] where stations share access to a mesh formed by a set of overlapping rings. Media access is performed by having a station capture one out of a number of tokens circulating within the network. Couplings between the rings are implemented by the user stations in a very simple fashion and under token control. As in a token ring, the network supports the transmission of variable length packets. In the results presented, a dual token ring is used as a basis for performance comparisons with the token grid. This is because unlike other proposed networks, the station transmission hardware requirements for the token grid are identical to that of a dual ring network such as FDDI [4], [3]. Each station in both networks has exactly two transmitters and two receivers. Also, both networks may be implemented with a very small station latency where only a few bits per station need be buffered. Thus it is expected that the token grid design may be very cost-effective. In addition, the total throughput of the token grid may be many factors higher than that associated with a conventional dual ring or bus network [4], [9].

The paper is organized as follows. Section II provides an introduction to the network topology. This is followed by a discussion of the basic token grid protocol in Section III. In Section IV the performance of the token grid is considered. We include a brief discussion of network maintenance and reliability in Section V. In Section VI some concluding remarks are made.

## II. THE TOKEN GRID NETWORK

The token grid is a two dimensional network structure arranged in  $R$  rows and  $C$  columns. An example for  $R = C = 4$  is shown in Fig. 1. Each station is identified by its row/column index as shown in the diagram. Also, associated with each row and column is a unidirectional ring which passes through all stations in the given row or column. Note that the same network topology has been used in a bridged-ring design [8]. The rings will be identified by their respective row or column IDs. For example, the top row ring is referred to as row-ring 0 and is abbreviated by RR0. Similarly, the right-most column ring in Fig. 1 is identified as column-ring 3 or CR3. Note that to simplify the description of the network, all of the row ring directions point to the right and the column rings point towards the bottom of the figure. The direction

1063-6692/94\$04.00 © 1994 IEEE

ACTIVISION, EA, TAKE-TWO, 2K, ROCKSTAR, Ex. 1019, p. 1 of 9

**EXHIBIT**

**Ex. 1019**






Fig. 1. The token grid (4 × 4 example).

of ring rotation is arbitrary and all subsequent discussion applies for Manhattan Street [6] reference directions also. It should also be noted that arbitrary numbers of stations can be accommodated simply by growing the network one row (or column) at a time. The new ring need not be fully populated. In this expository paper however, we consider only square networks (i.e.,  $R = C$ ).

In an implementation of the token grid, Fig. 1 represents an example of the virtual topology. The actual physical topology would more likely consist of a wiring concentrator with station fibres terminating at one or more centralized hubs [2], [4]. At the hub, passive connections are made between the fibres. It can be seen from Fig. 1 that each station in the token grid has both two transmitters and two receivers. This is the same as a dual ring network such as FDDI [4]. Thus in a wiring concentrator layout, the total amount of transmission hardware and fiber is identical for the two networks. Note that as in a typical ring, no transit packet buffering is provided in the stations except for the usual few bits of station latency.

Access to the various rings is determined by tokens which circulate throughout the network. Before an exact mechanism is discussed, we first consider the basic station configurations shown in Fig. 2. Note that the two configurations may be implemented by a pair of  $2 \times 1$  selectors at the station. Fig. 2(a) shows the double-ring (DR) connection. When in this configuration, the rings passing through the station are decoupled and it has access to each one independently. Fig. 3 shows a token grid where all stations are in the DR configuration. In this case, independent communication is possible on all of the row and column rings. To transmit on a given row or column ring, a row-token (RT) or column-token (CT) must be seized in a manner identical to the operation of a conventional token ring [2,4].

When a station is in the single-ring (SR) configuration of Fig. 2(b), the corresponding row and column rings have been concatenated into a single row/column ring. This ring is referred to as either a row/column ring (RCR) or a column/row ring (CRR). When this happens, the two rings are said to have been *merged*. An example is shown in Fig. 4 where station (1,2) has merged row ring 1 and column ring 2 thus forming RCR(1,2) (or CRR(1,2)). When a row/column ring




Fig. 2. Station configurations.




Fig. 3. The token grid (all stations in DR state).

is formed, two special tokens control media access to it. The first is referred to as a row/column token (or RCT) and the second is a column/row token (or CRT). In Fig. 4, the RCT is captured by stations on row 1 that transmit to stations on column 2. Similarly the CRT is used by stations on column 2 to transmit to stations on row 1. A protocol is used to ensure that at any time there is only one token active on any merged or unmerged ring. However note that there may be a number of stations which are simultaneously merged, provided that they do not share any of the same rows or columns. This restriction is enforced by a two-dimensional token passing algorithm to be discussed in the next section. It should also be noted that extra bits are added to identify the tokens as being either RT, CT, RCT or CRT. In the row/column and column/row tokens, the connected column or row is also identified.

### III. PROTOCOL OPERATION

A basic version of the token passing mechanism is now introduced. As discussed above, each row and column ring






Fig. 4. The token grid (station (1,2) in SR state).

has an associated token. In Fig. 1, for example, station (1,2) could access row ring 1 (RR1) by capturing row-token 1 (RT1) and column ring 2 (CR2) by capturing column-token 2 (CT2). Transmission to a station on a different row and column can only be accomplished when a ring merge has occurred. The merging process is identical for all stations and is now described. A formal description of the basic algorithm is shown in Protocol 1.

Consider a column token  $j$  arriving to a particular station  $(i,j)$ . When this happens, the token is always captured even though the station may have nothing to transmit. In the event that the station does have a packet for column ring  $j$ , it is permitted to transmit one packet at this time. The station continues to hold the column token until the row token appears. These actions are shown in the first section of Protocol 1.

When a row token arrives to a station which is not holding a column token, the station is free to capture it and transmit onto the ring in question. Afterwards, the token is released in the usual way.

When a row token  $i$  arrives at a station  $(i,j)$  which is currently holding a column token, the row token is always captured. At this time, the station may transmit a packet onto row ring  $i$  if desired. The station is now holding both row and column tokens and switches to the SR configuration thus merging its two rings. This is shown in the second section of Protocol 1. The station then issues a row/column token onto the merged ring through its row ring output fibre. Stations on this row are now able to capture this token and transmit to stations on the connected column. Note that this mechanism is similar to that used in the bridge design discussed in [7]. Any non-merged station that sees a row/column token is free to capture it and transmit onto the merged ring. Following this, the token is released.

Once the row/column token returns to station  $(i,j)$ , it is captured. At this time, the merging station is free to transmit a packet onto the merged ring. This action is shown in section 3 of the algorithm description. The token is then converted into a column/row token and circulated onto the station's column output fibre. This allows stations on the column to transmit packets to the connected row. As in the row/column case, any

#### PROTOCOL 1: BASIC TOKEN GRID PROTOCOL

##### Column Token $j$ Arrives To Station $(i,j)$ :

1. Capture the token.
2. If a packet is available, transmit onto Column Ring  $j$ .
3. Hold Column Token  $j$ . Wait for Row Token  $i$ .
4. Done.

##### Row Token $i$ Arrives To Station $(i,j)$ :

1. Capture the token if a packet is available for Row Ring  $i$  or if Column Token  $j$  is being held.
2. Transmit onto Row Ring  $i$  if a packet is available.
3. If the station is not holding Column Token  $j$  then release Row Token  $i$  and go to 5.
4. Enter the SR state. Generate and release Row/Column Token  $i,j$ .
5. Done.

##### Row/Column Token $i,j$ Arrives To Station $(i,k)$ :

1. Capture the token if a packet is available for Row/Column Ring  $i,j$  or if  $k = j$ .
2. If a packet is available then transmit onto Row/Column Ring  $i,j$ .
3. If  $k \neq j$  then release the token and go to 5.
4. Generate and release Column/Row Token  $i,j$ .
5. Done.

##### Column/Row Token $i,j$ Arrives To Station $(k,j)$ :

1. Capture the token if a packet is available for Column/Row Ring  $i,j$  or if  $k = i$ .
2. If a packet is available then transmit onto Column/Row Ring  $i,j$ .
3. If  $k \neq i$  then release the token and go to 5.
4. Enter the DR state. Regenerate and release Column Token  $j$  and Row Token  $i$ .
5. Done.

non-merged station that sees a column/row token is free to capture it and transmit. The token is released afterwards.

Eventually the column/row token will return to the merging station. At this time the token is captured and the station may again transmit a packet onto the merged ring. Following this, the station switches to the DR state and releases the original row and column tokens onto their respective rings. This is shown in the last section of the protocol description.

Using the algorithm described above, full network connectivity is achieved and governed by the rotation of the row and column tokens. However, it is apparent that transmissions on a given row or column ring do not necessarily occur in a contiguous fashion, that is, with the ring in an unmerged state for the duration of all transmitted column or row ring packets. Instead, individual transmissions may be interspersed with packets transmitted during ring mergings. In the resulting network, CT's rotate about each column, generating ring merges. Similarly, at each row, CT's are serviced by the row token, by creating merges in a fair and round-robin fashion.

In addition, it can be seen that if a station is permitted to capture a RT whenever it appears, then row-to-row traffic may be preferred over all others since column tokens must wait as part of the merging procedure. This advantage may be easily overcome by either restricting all transmissions to RCT and CRT captures, or by permitting only one RT capture each time a CT is observed. It is clear that this restriction may unnecessarily reduce the total throughput in certain nonuniform load situations. For example, consider the case where  $n$  stations on a given column are the only ones active, with packets to be sent on that column only. The above restriction will prevent any of these stations from using the full



bandwidth of its row ring. If a given station is using only a fraction  $1/n$  of the available column ring capacity, then at most it can achieve the same throughput on its row ring. Although the system is fair in this case, the total throughput has been reduced over that which is possible without this restriction.

In the scheme described above, a station passively waits for the appropriate token in order to transmit to a particular destination. However, it is pointless to merge a pair of rings if there are no stations on them wishing to communicate. A more efficient mechanism which also reduces delay under light loading may be used in this case. Each station is equipped with two state variables known as the Column Reservation Variable or CRV and the Row Reservation Variable or RRV. When CRV (or RRV) is set at a station, it indicates that some station on the associated row (or column) ring would like to acquire the column (or row) ring of the station. For example, if station (1,0) in Fig. 1 has a packet queued for station (3,2), it would set the CRV in station (1,2). This information is used by station (1,2) to determine whether a RCT (or CRT) is needed at the next merge opportunity. When a column token arrives at a station, the station checks both CRV and RRV. If either of CRV or RRV are set, then the station holds the token and waits for the arrival of the row token. When the row token arrives, the station generates either a CRT (if CRV=0 and RRV=1) or a RCT (if CRV=1 and RRV=0). Note that if CRV=RRV=1, the station generates a RCT first and then a CRT after the RCT returns.

The mechanism whereby a station sets the CRV (or RRV) bit in another station is very simple. Each row (column) token has a reservation flag bit for each column (row). Whenever a row (column) token passes a station, it checks to see if it has packets queued for any other column (row). If it does and has not yet made a reservation, it sets the corresponding bit(s) for the destination column (row) associated with the packets to be transmitted. When this token arrives at a station, the reservation flag bit for that column (row) is inspected. If the flag is set, then the station sets CRV=1 (RRV=1) and resets the reservation flag in the token. The same reservation flags are also included in the header of each packet. This allows a station to set the CRV (RRV) bit in the station at the beginning of each packet transmission.

It can be seen that transmission activity in the token grid has properties unlike conventional LAN's. As in other mesh networks, contention for various sets of links (i.e., rings) will be different depending upon destination traffic patterns. Naturally, stations contending for more bottlenecked rings will experience higher delays than others. Such a situation is unlike that in LAN's where a single channel must be shared by all transmitted packets, regardless of destination flows. It will be seen that because of the added transmission concurrency, station performance is much better in many common situations.

Throughout this paper, it will be assumed that a station may transmit at most one packet whenever a token is captured. In addition, the early token release mechanism used in FDDI [4] is employed. Thus, whenever a token is to be released following a packet transmission, it is done immediately after the transmission is completed. Versions using delayed token




Fig. 5. Uniform load capacity.

release [2] are also possible, but are much more sensitive to propagation delay [10].

#### IV. NETWORK PERFORMANCE

In this section, we discuss the performance of the token grid network. As a benchmark, the performance of a dual ring is used for comparison. In a dual ring (such as FDDI [4]), the transmission hardware requirements are identical to that of the token grid. Each station is equipped with two transmitters and two receivers. In addition, each station in both networks may be implemented with a very short station latency. In order to ensure a valid comparison, it is assumed that only one packet may be transmitted per token seizure in both networks. Since the token grid topology allows concurrent transmission of packets on disjoint routes, it is assumed that in the station implementation, separate packet queues are maintained for each column and row. This simple function would typically be implemented by a set of pointers in a low level device driver or VLSI DMA controller. For comparison purposes, we assume a wiring hub layout for both networks. It is assumed that the distance from each station to the wiring center is the same for all stations. The token ring consists of  $N$  stations with a total ring latency defined to be  $\tau$  seconds. The station-to-station latency is thus  $\tau/N$  seconds. In the results, our attention is restricted to square token grids with  $R = C$  and a total of  $N = R \cdot C$  stations.

The maximum throughput or capacity performance of the token grid is considered first. Since the token passing mechanism yields complex cycle patterns, exact simulation results for network capacity are first given. Subsequently, a simple capacity approximation will be described.

In Fig. 5, the normalized network capacity is shown plotted against the normalized total ring latency  $a = \tau/T$  where  $T$  is the packet transmission time in seconds. Results are given



for  $4 \times 4$ ,  $6 \times 6$  and  $10 \times 10$  networks. All stations were assumed to be permanently backlogged and transmitting fixed-length packets into the network uniformly distributed across all destinations. The capacities of the corresponding 16, 36 and 100 station dual token rings are also included in the graphs and are given by

$$C_{DTR} = \frac{2}{1 + a/N + t_{\text{token}}/T} \quad (1)$$

where  $t_{\text{token}}$  is the token transmission time. Equation (1) is easily obtained by observing that at capacity, the time between successive packet transmissions is given by  $T + t_{\text{token}} + \tau/N$  seconds. Note that the factor of 2 results because of the two rings. In the curves it was assumed that  $t_{\text{token}}/T = 0.1$ .

In Fig. 5 the capacity for the token rings drops off from a value slightly less than 2 in the usual fashion. However, the total throughput for the token grid is much higher. In the 16 station example, the capacity is more than 100% higher for small values of  $a$ . Also it is important to note that unlike conventional LAN's [4], the total capacity of the token grid is an increasing function of the network size. In the 100 station case, the total throughput is more than 5 times that of the token ring at small values of  $a$ . These improvements result from the transmission concurrency obtained by using a network topology which supports multiple concurrent transmissions.

Under the conditions shown, a simple and accurate approximation is given for the uniform traffic capacity. The approximation is obtained by assuming that network activity is cyclic and that the column tokens are never idle. To generate this behaviour, assume that the packet lengths are fixed and that all row/column transmission modes are exactly synchronized. In Fig. 1 for example, assume that at some time instant stations (3,0), (2,1), (1,2) and (0,3) have all created row/column tokens. Since all stations are busy, transmission activities and timing on the 4 row/column rings will be identical and a series of simultaneous transmissions occur. After the RCT and CRT rounds are completed, each row and column token are also used by the following stations at the same time. This results in 8 concurrent row and column ring transmissions followed by a new set of four ring mergings. We denote a cycle to be the time between these merging instants. It can be seen that in this example, four cycles identical in duration occur before the same set of stations merge again. Employing a usual cycle analysis, the capacity can be shown to be

$$C_{TG} \approx \frac{\sqrt{N} + 1}{(1 + t_{\text{token}}/T)(1 + \frac{1}{2\sqrt{N}}) + \frac{a}{2N}(2 + 1/\sqrt{N})}, \quad (2)$$

or

$$C_{TG} \approx \frac{\sqrt{N}(\sqrt{N} + 1)}{2\sqrt{N} + 1} C_{DTR} \quad (3)$$

In deriving the above equation, the total normalized throughput per cycle is given by  $U = R \cdot 2(R + 1)T$ . In this cycle,  $R$  stations simultaneously release row/column tokens. For each of these concurrent cycles, there are a total of  $2R + 2$  packets transmitted. One is associated with each of the row and column tokens, and the other  $2R$  are for the stations

which transmit using the RCT and CRT tokens. This gives the value for  $U$ . To obtain the final result,  $U$  is divided by the average time per cycle  $L$ . This is easily shown to be  $L = T + 2R(T + \tau/R^2 + t_{\text{token}}) + t_{\text{token}} + \tau/R^2$ . In Fig. 5, results are shown for this capacity approximation. As expected, it tends to slightly overestimate the true capacity since in a real system, column tokens would have a nonzero waiting time. It can be seen however that the approximation is a very good one. Also, it can be seen that when  $N$  is large,  $C_{TG} \approx \frac{\sqrt{N}}{2} C_{DTR}$ .

It is also interesting to consider a comparison of the token grid to a dual ring under light load conditions. We will assume that the merge reservation mechanism discussed in Section III is used and that tokens are received entirely before being relayed. In this case the worst-case delay for a square token grid is easily shown to be  $(5R - 2)(t_{\text{token}} + \tau/N) + T + (2R - 1)\tau/N$ . The corresponding expression for a dual ring is  $R^2(t_{\text{token}} + \tau/N) + T + R^2\tau/N$ . Thus the light load delay is proportional to  $N$  for a dual ring and to  $\sqrt{N}$  for the token grid. A similar analysis for mean delay under light load results in the same conclusion. The above result shows that when  $R > 4$ , the token grid delay is always better. When this is not the case, the dual ring delay may be better depending upon the relative magnitudes of  $t_{\text{token}}$  and  $\tau$ .

#### A. Mean Delay Bound

In this section, a lower bound is derived for the mean waiting time of packets under uniform symmetric traffic conditions. In the token grid protocol, it is readily apparent that there are many variations in the how stations transmit packets to a given destination. For example, station (0,0) can transmit to station (0,3) upon capturing either row token 0 or any of the RC tokens generated by stations on row 0. To simplify the analysis, we will assume that all traffic in the network is transmitted using only RC and CR token captures. Stations thus transmit to others on the same row or column after capturing the appropriate RC and CR tokens generated by other stations on the same row or column. Accordingly, each station maintains a set of  $2(R - 1)$  queues, corresponding to the identities of the RC and CR tokens it captures for transmission.

Consider a cycle which consists of the total excursion time of a particular column token around its column. For simplicity, we will suppress all time indices and view the rotation of the token in cycle 0. Also for convenience we will consider the column 0 token (CT0) starting its cycle at station (0,0). Note that this cycle comprises  $R$  subcycles each consisting of the time spent at each of the  $R$  rows before returning to row 0. At any particular row  $r$ , the time delay incurred by the column token is given by the sum of three terms, namely

$$W_{RT}^r + T_{RC}^r + T_{CR}^r \quad (4)$$

where  $r \in \{0, \dots, R - 1\}$  where

$W_{RT}^r$  = waiting time for the row token  $r$  to arrive at station (r,0)

$T_{RC}^r$  = total duration of the RC token rotation cycle at row  $r$

$T_{CR}^r$  = total duration of the CR token rotation cycle at row  $r$

The two token rotation cycles are associated with the merging of station (r,0) once the row token arrives. Note that



95%  
Sep  
resu  
stati  
bou  
whe  
muc  
It  
unde  
bou  
Upo  
load  
max  
a co  
requ  
later  
further

The  
as a  
Thes  
can  
espe

B. N

In  
toke  
is r  
inter  
seco  
mod  
In  
a se  
traff  
serv  
stati  
from  
all c  
Res  
stati  
assu  
whe  
tran  
Tok  
stati  
In  
Sep  
both  
the  
In t  
tota  
han  
perf  
thro  
and

$W_{RT}^r$  is exceedingly complex and dependent upon the global coupling between row and column tokens.

When the RC and CR tokens are rotated by station  $(r, 0)$ , a total of  $2(R-1)$  stations are given the opportunity to transmit. As a result, over the complete column token cycle, stations on column 0 are "polled" a total of  $R-1$  times (once for each CR token generated by other stations on the column). However, all other stations are polled only once. This complication is handled by noting that each of these polling actions addresses a different queue at each station. It can be seen that over the complete column token cycle, a total of  $2R(R-1)$  distinct queues are polled. We will thus view this column token cycle as a cyclic server system consisting of  $2R(R-1)$  "stations" each containing a single queue. As in [11], we will refer to the time between polling instances as the "walk-time" of the token.

Since stations transmit all packets via row/column and column/row tokens, a merging station never transmits upon the arrival of the RC or CR tokens that it generates. As a result, even if all links are of equal length, the walk-times are not the same for all stations. It can be seen that there are three different walk-times as follows

$$\begin{aligned} w_1 &= \tau/N + t_{token} \\ w_2 &= 2(\tau/N + t_{token}) \\ w_3^r &= 3(\tau/N + t_{token}) + W_{RT}^{r+1}. \end{aligned} \quad (5)$$

Here,  $\tau$  is taken modulo  $R$ . Queues with walk-times given by  $w_1$  correspond to those passing an RC or CR token to another non-merged station. When this happens, the following station will transmit if the appropriate queue is nonempty. Queues with walk-times of  $w_2$  are those who pass the RC token back to the merging station that generated it. Since this station will not use the token to transmit, it immediately releases a CR token onto its column. As a result, there is a walk-time of twice  $w_1$  before the next queue is polled. Similarly, those stations passing the CR token back to the merging station experience a walk-time of  $w_3^r$ . In this case, the token must first proceed back to the merging station. The column token is then regenerated and passed to the next station on the column. Upon receiving the column token, the station must wait for the row token. Once the RT arrives, the station merges and releases the RC token. This accounts for the final propagation and token transmission terms. In each column token rotation cycle, there are  $2R(R-2)$  queues which experience  $w_1$ . Half of the remaining  $2R$  queues have walk-times of  $w_2$  and the other half  $w_3^r$ .

In addition to all previous assumptions, we will also assume that each of the  $2R(R-1)$  queues has unlimited storage and is fed by an independent Poisson arrival process with identical mean arrival rate  $\lambda_q$ . From the above discussion, we can see that the total walk-time over the entire cycle is given by

$$\begin{aligned} W_t &= 2R(R-2)w_1 + R w_2 + \sum_{i=0}^{R-1} w_3^r \\ &= 2R(R-2)(\tau/N + t_{token}) + 2R(\tau/N + t_{token}) \\ &\quad + 3R(\tau/N + t_{token}) + \sum_{i=0}^{R-1} W_{RT}^i, \end{aligned} \quad (6)$$




Fig. 6. (a) and (b) Mean delay bound, 16 and 36 stations.

or

$$W_t = R(2R+1)(\tau/N + t_{token}) + \sum_{i=0}^{R-1} W_{RT}^i. \quad (7)$$

To obtain a lower bound on the mean waiting time experienced by packet arrivals, we will assume that  $W_{RT}^i = 0$  for all  $i$ . As a result, we now have a limited-service cyclic polling system with fixed (but non-equal) walk-times. Each of the  $2R(R-1)$  polled queues correspond to a station in a conventional polling system [11]. In view of the fact that the walk-times have zero variance, a closed-form expression is available for the mean waiting time [12],

$$\bar{W}_b = \frac{(1 + \rho_0)W_t + \rho_0 T}{2(1 - \rho_0 - \lambda_q W_t)} \quad (8)$$

where  $\rho_0 = 2R(R-1)\lambda_q T$ .

From the symmetry imposed on the system, the overall mean waiting time of the system is also given by this expression. Since we have assumed that the row token waiting times are zero, the expression for  $\bar{W}_b$  gives a strict lower bound on the mean waiting time which could be experienced in the network.

Simulation experiments were performed to compare the lower bound with the exact mean packet waiting time. The results are shown in Fig. 6(a) and (b). In both figures the lower bound is shown as the dashed curves. The simulation results are given by the unconnected points. In all cases, the



95% confidence intervals are within 10% of the plotted values. Separate curves are given for various values of  $a$ . In Fig. 6(a) results for a 16 station network are given and in Fig. 6(b) a 36 station token grid was used. It can be seen that in all cases the bound closely predicts the capacity of the real network. Also, when the propagation delay is small to modest, the bound is much more tight.

It is also possible to estimate the error in the bound under light load conditions. As propagation delay grows, the bounding assumption that  $W_{RT}^i = 0$  is increasingly inaccurate. Upon completion of a set of concurrent mergings under light loading, a row token must travel a minimum of 1 hop and a maximum of  $\sqrt{N} - 1$  hops, until it reaches the station holding a column token. To estimate the error, we assume that  $W_{RT}^i$  requires approximately one half the total row ring propagation latency on average. Since there are  $R$  of these terms, and further assuming that an arriving packet must on average wait for half of this additional cycle time, the error prediction is

$$\bar{W}_e = \frac{R^2}{4}(\tau/N + t_{\text{token}}). \quad (9)$$

The sum of the error and bound (i.e.,  $\bar{W}_b + \bar{W}_e$ ) may be used as a better mean delay approximation than the bound itself. These curves have also been included in Fig. 6(a) and (b). It can be seen that this results in a very good approximation, especially under low to moderate propagation delay.

### B. Nonuniform Load Experiments

In this section, results are presented which subject the token grid to different nonuniform load examples. The first is referred to as a "server model" and models the traffic interaction between a set of server and non-server stations. The second shows an example of a community-of-interest traffic model.

In the server model, each station is assumed to be either a server or nonserver station. Each nonserver station sources traffic at a fixed rate  $\lambda_{ns}$  to each server station. Similarly, each server station sources traffic at the same rate to each non-server station. In the case where there is a single server, all traffic flow from the nonserver stations focuses on the server. Similarly, all other traffic is sourced from the server to all other stations. Results are given in the form of mean delay vs. non-server station throughput as the latter is varied. In the graphs we assume a total of 36 stations in a wiring concentrator layout where each station is located 200m from the hub. The channel transmission rate is 100 Mbps with 1000 bit data packets. Token packets are assumed to be 100 bits. Arrivals to each station occur via a Poisson Process.

In Fig. 7(a) results are given for a network with two servers. Separate curves are given for the mean delay incurred by both server and nonserver stations. Also results are given for the same traffic scenario when a dual-ring network is used. In both cases, the throughput of the server stations limit the total capacity of the network. This is expected since they are handling most of the network traffic. However, the overall performance of the token grid is superior. In this case the total throughput is roughly twice that of the dual ring. In Fig. 7(b) and (c), similar results are given for a network with six and




Fig. 7. (a), (b), and (c) Server model, 2, 6, and 12 servers.

twelve servers. Again it is seen that the token grid is much better. In the six and twelve server cases the advantage in total throughput increases over that in the two server case.

A single set of curves is also presented for the community-of-interest (COI) traffic model. This is given in Fig. 8. In this set of experiments, the stations were partitioning into COIs. When an arrival occurs to a station in one COI, then with probability  $P_{COI}$  a destination within that COI is uniformly selected. Similarly, with probability  $1 - P_{COI}$  a destination outside its COI is uniformly chosen. By varying  $P_{COI}$ , the level of interaction between COI's can be adjusted. To simplify the graphs, we only include the extreme cases corresponding to  $P_{COI} = \{0, 1\}$ . We assume the same physical parameters as that in the server model experiment. In all cases it was






Fig. 8. Community-of-interest model.

assumed that the COI's were known in advance so that a reasonable physical partitioning could be used. In practice this would be accomplished by changing the connections made between stations at the hub. In the 6 and 12 COI results, the partitioning was determined by assigning COI's from left to right starting at the top of the network. In the 12 RC COI case, the COI's were formed by ensuring that each COI had sole use of a row or column ring. In this case, when  $P_{COI} = 1$  a very large degree of concurrency is possible. In all cases it can be seen that even when the community traffic is not heavy, much better performance is attained than over a dual ring. In both Figs. 7 and 8, the 95% confidence intervals are within 15% of all plotted points.

In the extreme case, if stations can be placed so that all communications are between stations in the same row or column, then the maximum throughput configuration is achieved. In this case, the token grid behaves as a set of  $R+C$  decoupled rings. Using a standard cycle length procedure, the total capacity (for  $R = C$ ) is given by  $C_{coi} = 2\sqrt{N}/(1 + t_{token}/T + a/N)$ . This is  $R$  times the corresponding value for a dual token ring under the same assumptions. It is interesting to note however that in certain focused load situations a dual ring may have higher throughput than the token grid. This is a product of the topology and token passing mechanism used. For example, in the token grid two stations on the same row ring can at best communicate between each other at a maximum throughput equal to the link rate. However, the same two stations in a dual ring could conceivably use both rings to communicate. The same two stations connected to different rows and columns in the token grid could support the same mutual throughput as the dual ring. This observation is also true when comparing a dual bus such as IEEE 802.6 [9] with a dual ring. Since there is only a single path from any given station to another, one pair of stations can at most achieve half the throughput of the dual ring case. The ability for a network to support a much higher load when a large number of stations are active is a much more important factor.

## V. MAINTENANCE AND RELIABILITY

In this section a brief discussion is given of the maintenance and fault tolerance of the token grid. Initialization and maintenance of the tokens is discussed first. Since each station

is connected to two independent rings, procedures similar to those in IEEE 802.5 [2] and FDDI [4] may be adopted. Following the technique in [2], a ring monitor is established for each row and column ring. As in the token ring case, this may easily be done using a distributed competition based upon a "token claim" procedure. A station that successfully establishes itself as a monitor generates the appropriate row or column token. An activity timer is used to initiate error recovery following the loss of these tokens. Each station is also equipped with a merge activity timer which is initiated whenever the station merges its rings. The expiry of this timer indicates the loss of the RCT or CRT and the station merely switches back to the DR state. Token recovery is then provided by the monitor stations. Other token maintenance functions such as checking for eternal frames and token duplications follows directly from the token ring case [4].

The issue of reliability in the presence of station power and link failures is now discussed. In the LAN environment, network sensitivity to station power failure is a much more critical factor since physical link disruptions are much less likely. The total performance of the token grid is reasonably insensitive to catastrophic link failures and the network degrades gracefully in this case. A single link failure along a row or column ring will result in the loss of connectivity only between those stations. In an  $R \times R$  network, this represents a loss in connectivity of a fraction  $1/(R+1)$  of the full network connectivity. Thus in a 100 node network, roughly 1% of the full connectivity would be lost due to a single physical link outage. Similarly, two link failures can at most disrupt 2 of the  $2N$  rings in the network. The remaining  $2(N-1)$  rings would be still operational thus permitting communications amongst a large subset of the stations. This is due to the fact that both row/column and column/row paths are being used. There are a number of options with regard to operation in the presence of station power failures. As in FDDI [4], it is assumed that in the physical implementation, optical bypass switches are used to maintain ring integrity following a station power failure. When a power failure occurs, the station may be either bypassed in the SR or DR configuration. In [10], a simple method has been devised which maintains full network connectivity regardless of the number of failed stations. In this technique, certain stations are designated to fail in a particular switching configuration. As a result, the network eventually degrades into a worse-case scenario where all remaining stations are connected to a single ring running at the channel data rate. A full discussion of this feature is lengthy and the reader is referred to [10].

## VI. CONCLUSIONS

In this paper the token grid network concept was introduced. Unlike conventional LAN's and MAN's, media access in the token grid is performed over a two-dimensional grid topology. However, the stations may be realized with about the same hardware complexity and latency as in a token ring network. The token grid has a number of advantages over conventional approaches. Unlike typical shared media networks which have a fixed a priori upper bound on throughput, the maximum



throughput of the token grid under uniform load increases with the number of stations. In addition, because of the transmission concurrency and path diversity in the network, the maximum throughput under uniform traffic conditions may be many times that of a dual token ring with the same transmission hardware.

#### ACKNOWLEDGMENT

The author would like to thank E. Hahne for many useful discussions and for her comments on an earlier version of this paper. Thanks is also extended to the reviewers for a number of excellent suggestions.

#### REFERENCES

- [1] ANSI/IEEE Std. 802.3, "Carrier sense multiple access with collision detection," Institute of Electrical and Electronics Engineers Inc., 1985.
- [2] ANSI/IEEE Standard 802.5, "Token ring access method and physical layer specifications," 1985.
- [3] F.E. Ross, "An Overview of FDDI: The Fiber Distributed Data Interface," *IEEE J. Select. Areas Commun.*, vol. 7, pp. 1043-1051, Sept. 1989.
- [4] American National Standard, "FDDI token ring media access control (MAC)," ANSI X3.139, 1987.
- [5] A.S. Acampora, M.J. Karol, and M.G. Hluchyj, "Multihop lightwave networks: A new approach to achieve terabit capabilities," in *Proc. ICC'88*, vol. 1, pp. 1478-1484, 1988.
- [6] N.F. Maxemchuk, "Regular and mesh topologies in local and metropolitan area networks," *AT&T Tech. J.*, vol. 64, pp. 1659-1686, Sept. 1985.

- [7] T.-M. Tsai and L. Merakos, "Interconnection of High Speed Token Ring LANs: A Switch-Connection Approach," in *Proc. INFOCOM'90*, San Francisco, CA, 1990, pp. 989-996.
- [8] S. V. Kartalopoulos, "A Manhattan fiber distributed data interface architecture," in *Proc. GLOBECOM'90*, San Diego, CA, Dec. 1990, pp. 141-145.
- [9] "Distributed dual bus (DQDB) subnetwork of a metropolitan area network (MAN)" IEEE 802.6 Proposal and submissions, Feb. 7, 1990.
- [10] T.D. Todd and E.L. Hahne, "Local and Metropolitan MultiMesh Networks," in *Proc. ICC'92*, Chicago, IL.
- [11] H. Takagi, *Analysis of Polling Systems*. Cambridge, MA: The MIT Press, 1986.
- [12] K. S. Watson, "Performance evaluation of cyclic service strategies - A survey," *Performance'84*, North-Holland, Elsevier Science, 1984, pp. 521-533.



**Terence D. Todd** (S'78-M'79-M'84) received the B.A.Sc., M.A.Sc., and Ph.D. degrees in electrical engineering from the University of Waterloo, Waterloo, Ont., Canada, in 1978, 1980, and 1984, respectively.

During that time he spent 3 years as a research associate with the Computer Communications Networks Group at the University of Waterloo. He is currently an Associate Professor of Electrical and Computer Engineering at McMaster University in Hamilton, Ont., Canada. He spent 1991 on research leave in the Distributed Systems Research Department at AT&T Bell Laboratories, Murray Hill, NJ. He is currently Head of the Optical Network Protocols project in the Telecommunications Research Institute of Ontario. His current research interests include Metropolitan/Local area networks and the performance analysis of computer communication networks and systems.