Performance Evaluation of the MetaRing MAC Protocol Carrying Asynchronous Traffic

G. Anastasi, L. Lenzini

University of Pisa, Department of Information Engineering
Via Diotisalvi 2 - 56126 Pisa, Italy
anastasi@iet.unipi.it, lenzini@iet.unipi.it

ABSTRACT

This paper analyzes in depth the asynchronous part of the MetaRing MAC protocol by considering a network configuration (hereafter scenario) in which the slot reuse capability is used to the minimum extent possible. The MetaRing MAC protocol behavior is first analyzed in *asymptotic conditions* (i.e., when each station is trying to seize all the medium capacity), then in *underload conditions* (i.e., when the offered load is lower than the maximum achievable aggregate throughput), and finally in *overload conditions* (i.e., when the offered load is slightly higher than the maximum achievable aggregate throughput). Our results show that the MetaRing is unfair and that the unfairness depends, in particular, upon the workload parameters, in addition to the K and L protocol parameter values. Our analysis includes medium capacities in the range 150 Mbps - 2.4 Gbps.

1. Introduction

The MetaRing MAC protocol has been analysed in ([CIDO93],[WU92]). In both papers the cell destination addresses are assumed to be uniformly distributed among half of the ring. Results show that the MetaRing MAC protocol, with a uniform distribution for destination addresses, provides all the stations with the same opportunity to access the medium.

A more accurate performance analysis of the MetaRing MAC protocol should take into consideration more general traffic patterns and station configurations. Due to the complexity of the MAC protocol, this is clearly a very difficult task.

Therefore, to further advance the MetaRing MAC protocol analysis carried out in ([CIDO93],[WU92]), we focus on a scenario [ANAS95] which belongs to the class of worst-case scenarios, which is relatively simple to analyze, but still provides useful information on the MetaRing performances. This scenario consists of a group of N+1 equally spaced stations located in one half of the ring as shown in Figure 1.

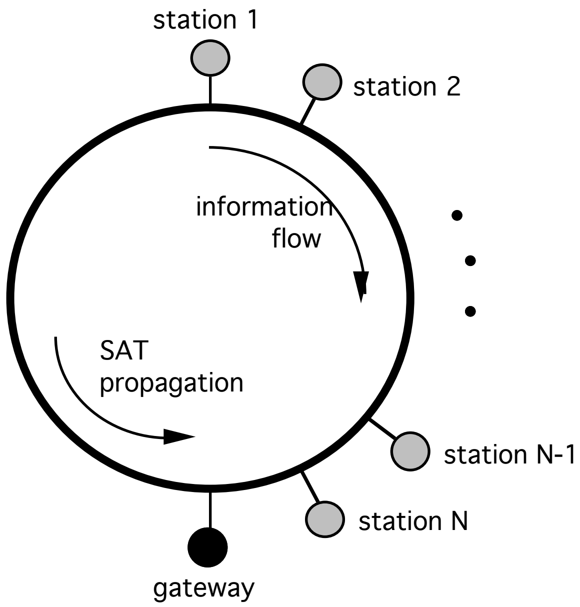

Figure 1: Network scenario

The $(N+1)-th$ station operates as a gateway interconnecting the MetaRing with other subnetworks. Furthermore, it is assumed that stations from $\{1\}$ to $\{N\}$ transmit asynchronous traffic to the gateway. Since slots are emptied only by the gateway, a slot can be utilized at most once during its journey from station $\{1\}$ to the gateway. Also, stations observe different patterns of empty/free slots. Specifically, station $\{1\}$ always observes empty slots, while station $\{N\}$ is covered by traffic transmitted by its upstream stations and therefore it can send a cell only if a slot has been left empty by stations $\{1,2,...,N-1\}$.

It can be easily verified that this scenario is particularly critical with respect to slot reuse since the spatial reuse capability is used to the minimum extent possible. Also, stations observe different patterns of empty/free slots.

Performance measures of interest, such as average access delay, throughput, and packet loss are then estimated. The MetaRing MAC protocol behavior is first analyzed in asymptotic conditions, and then in underload and overload conditions.

Specifically, under asymptotic traffic conditions a closed formula for the throughput achieved by any station is derived, while in underload and overload conditions a simulative analysis is performed. The rationale behind the simulative choice is that the modeling and performance analysis of the MetaRing MAC protocol is known to be a very difficult problem. The reason for this is the high degree of interactions among a plethora of processes, which makes an exact analysis of the network almost impossible.

This paper is organized as follows. Section 2 describes the MetaRing MAC protocol. Section 3 outlines the analysis in asymptotic condition. Section 4 is devoted to the analysis in underload conditions, while Section 5 reports the results related to the analysis in overload conditions. Conclusions are drawn in Section 6.

2. MetaRing MAC Protocol Description

The MetaRing is a Medium Access Control (MAC) protocol proposed for LANs and MANs operating at a speed of 1 Gbit/sec and above. It connects a set of stations by means of a bi-directional ring made up of full duplex point-to-point serial links. A MetaRing can operate under two basic access control modes: *buffer insertion* for variable size packets, and *slotted* for fixed length cells. In this paper only the slotted access mode is considered.

The MetaRing MAC protocol provides two types of services: asynchronous and synchronous. This section outlines the MAC protocol aspects which are relevant to the study reported in the paper (i.e., asynchronous traffic only). Details on the MAC protocol can be found in [CIDO93], [WU92], and [OFEK94].

According to the slotted access mode, information is segmented into cells and each cell transmitted in one slot. Slots are structured into a header and an information field. The header includes a busy bit which indicates whether the slot is empty (busy bit=0) or busy. A cell can be accommodated on an empty slot. Since the ring is bi-directional the MetaRing MAC protocol uses the *shortest path* criterion to choose one of the possible directions. Cells are removed by the destination station which frees the slot by resetting the busy bit to zero. After cell removal the slot can be reused by the same station or by its downstream stations. Slot reuse increases the aggregate throughput beyond the capacity of single links but may cause starvation. This happens if some station is constantly being "covered" by an upstream ring traffic and thus is not able to access the ring for a very long period of time.

In order to prevent this, the MetaRing MAC protocol includes a fairness mechanism which is based on a control signal, called *SAT* (*SATisfied*), which we assume circulates in the opposite direction with respect to the flow of information it regulates. Each station has a counter which is increased by one every time the station transmits a cell. The station can transmit cells every time it observes an empty slot, unless the counter is equal to the value of a MAC protocol parameter denoted by K. When the counter reaches the value of K, the station must refrain from sending new cells until the *SAT* signal arrives at the station. Upon *SAT* arrival, the station resets the counter to zero if it has transmitted a number of cells greater than or equal to an additional protocol parameter denoted by L, with $K \geq L \geq 0$, and then releases the *SAT* towards its upstream station. When the counter is less than L the station is not satisfied, and the following two mutually exclusive events can occur:

- if the station has data to transmit, it holds the *SAT* until the counter reaches the value of L, then the station resets the counter and releases the *SAT*;
- if the station has no data to transmit, the station resets the counter and releases the *SAT*.

3. Analysis in Asymptotic Conditions

In this section we analyze the MetaRing behavior when every station tries to seize all the medium capacity (*never-empty station queues*). This is the situation that causes the maximum interference between stations, since all compete in grabbing as much medium bandwidth as possible. Although this situation should never occur in a real network, its investigation is very useful as it provides insight into the MAC protocol's capability to solve conflicts between stations fairly. Performance measures of interest in asymptotic conditions are the aggregate throughput (ρ) that the MetaRing can manage, and the throughput that each station can achieve on the ring $(\gamma_i,\ i = 1,2,\ldots,N)$. The ratio between the most and the least station throughputs achieved has been adopted as a metric for the MetaRing fairness in asymptotic conditions. Note that results obtained in this section are only applicable to the network configuration reported in Figure 1.

Hereafter, the throughput achieved by a station is expressed as the average number of slots made busy by that station in a ring latency interval, normalized to the number (S) of slots circulating in the ring (ring latency expressed in slots). Therefore, the aggregate throughput gives a measure of the number of times a slot is utilized by stations during a ring round. Thus, in our scenario, the aggregate throughput can never exceed 1.

In asymptotic conditions the MetaRing behavior is deterministic. If we define the *transmitting state* of a MetaRing at time *n* by the following vector $\left\{k^{(SAT)}, Count(1), Count(2), ..., Count(N)\right\}_{(n)}$ where $1 \leq k^{(SAT)} \leq S$ is the slot position (measured, say, from station $\{N\}$) of the *SAT* within the ring, and $Count(j),\ j = 1,2,..,N$, is the value of the counter at station $\{j\}$, we will prove that the MetaRing converges to a steady state operation, such that active stations visit a finite number of transmission states on successive steady state cycles, hereafter called *SAT cycles*. More formally, it will be proved that the process

$$Y = \left\{\left[k^{(SAT)}, Count(1), Count(2), ..., Count(N)\right]_{(n)}, n = 0,1,2,...\right\}$$

is regenerative with respect to the time instants (i.e., regeneration epochs) at which the *SAT* leaves station $\{N\}$ while this station is observing a stream of empty slots. Hence, the aggregate(station) throughput expression is derived by evaluating the fraction of time stations(station) transmit(s) in a *SAT* cycle (i.e., the time interval between two consecutive regeneration epochs). Without losing any generality, throughout the paper it is assumed that the *SAT* cycle is evaluated by observing the *SAT* rotation from station $\{N\}$. Trajectories of the *Y* process related to several system parameter values, are plotted below in several figures. Each figure plots, as a function of time, the *SAT* position along the ring in addition to the evolution of $Count(j),\ j = 1,2,....,N$ starting from the time station $\{j\}$ is hit by the *SAT*. The *SAT* position is represented by a straight line with a slope equal to +1. Each $Count(j),\ j = 1,2,....,N$ evolution is represented by a horizontal segment. There is one segment for each station. For ease of representation both the *SAT* position and the $Count(j)$ evolutions are represented by continuous lines. Stations are reported in the ordinate of each figure. Each horizontal segment is interrupted when either the station observes busy slots or when the station has transmitted K cells, whichever occurs first. In the former case the $Count(j),\ j = 1,2,....,N$ evolution is resumed when empty slots are again observed by the station, while in the latter situation the evolution restarts when the *SAT* is observed again during the next round around the ring. When a station has transmitted its maximum quota of cells (i.e., K cells) and before the *SAT* is observed again, each emptied slot observed by a station is left unchanged by the station itself. With this type of representation SAT cycles can be identified straightforwardly, as it is easy to realize when the process Y restarts visiting the same set of states already visited in the previous cycle.

In the following two sections we report the closed formulae for the aggregate and station throughputs as a function of the ring latency (S), the number of active stations (N), the distance between stations (d) which is assumed constant, and the protocol parameters K and L. Both S and d are expressed in slots and they are assumed to be integers greater than or equal to one.

3.1 Aggregate Throughput

Two alternatives need to be considered: $K \cdot N < S$ and $K \cdot N \geq S$. If $K \cdot N < S$, all stations transmit K cells and then stop, and wait for the SAT to come back. In this case the SAT cycle length is equal to the round trip S. Since for $K \cdot N < S$ there is a fraction of slots which is not utilized, it follows that $\rho < 1$. If, on the other hand, $K \cdot N \geq S$, all the slots within a SAT cycle are utilized for cell transmissions and hence $\rho = 1$. Thus

$$\rho = \min\left(\frac{N \cdot K}{S}, 1\right) \qquad [1]$$

Hence, with N and S fixed, the aggregate throughput linearly increases with K up to the value 1.

3.2 Station Throughputs

For the evaluation of the station throughputs several alternatives need to be considered. Let us start with $K \cdot N < S$ and $K \cdot N \geq S$.

CASE 1: $K \cdot N < S$

In this case the ring is clearly underutilized. If a single station is active $(N = 1)$ the throughput (γ) it achieves coincides with the aggregate throughput. Therefore, from [1] we derive

$$\gamma = K/S \qquad [2]$$

On the other hand, if $N > 1$ it can be readily verified that any station transmits K cells during each *SAT* rotation which lasts S slot time (each station releases the *SAT* just after its reception). Hence, the *SAT* cycle has a duration of S slots and the throughput achieved by station $\{i\}$ (for $i = 1,2,...,N$) is $\gamma_i = K/S$.

If $K \cdot N \geq S$ we need to make a further distinction depending on whether $K(N-1)+L \leq S$ or $K(N-1)+L > S$. For both cases, to derive the trajectory of the Y process it is convenient to start from the state where the *SAT* leaves station $\{N\}$ and $Count(i) = K$, for $i = N-1, N-2,...,2,1$. For ease of representation, throughout the text this state will be denoted by

$$\{\Phi\} = \left\{k^{(SAT)} = 1, Count(1) = K,......, Count(N-1) = K, Count(N) = 0\right\}$$

In [ANAS96] it is proved that $\{\mathbf{\Phi}\}$ is always visited by the Y process whatever the initial state of Y is, and furthermore, from the time $\{\mathbf{\Phi}\}$ is visited, station $\{N\}$ observes a continuous stream of $2d$ empty slots before being covered by the slots made busy by its upstream stations. Also, from $K \cdot N \geq S$ and $d = \lceil (S/2)/N \rceil$ (i.e., stations are equally spaced in half of the ring), the condition $K \geq 2d$ readily follows. This condition will be used extensively throughout this paper.

CASE 2: $K \cdot N \geq S$, $K \cdot (N-1) + L \leq S$ - *the SAT circulates along the ring without being withheld by station* $\{N\}$

Starting from the initial state $\{\mathbf{\Phi}\}$ (which occurs at epoch t_1 in Figure 2), by the time the *SAT* reaches station $\{N-1\}$ thus enabling it to transmit, station $\{N\}$ has transmitted d cells. Furthermore, from the time the first empty slot filled in by station $\{N-1\}$ reaches station $\{N\}$, station $\{N\}$ has transmitted another additional quota of d slots. Hence, $2d$ slots have been transmitted by station $\{N\}$ since the time station $\{N\}$ itself released the *SAT* before being covered by the busy slot string generated by the upstream stations. Bearing in mind the MetaRing *MAC* protocol operations, it is straightforward to verify that this process will continue for a generic station $\{j\}$, $j = 2,...,N$, which can thus transmit $2d$ cells from the time it forwards the *SAT* to station $\{j-1\}$ until the head of the busy slot string generated by station $\{j-1\}$ has covered station $\{j\}$ (note that $K \geq 2d$). On the other hand, since station $\{1\}$ observes only empty slots, it can transmit continuously K cells from the time the *SAT* passes in front of it. Station $\{1\}$ refrains from transmitting after K cell transmissions, and then allows slots to remain empty. Those slots propagate along the ring towards station $\{2\}$ which thus uses them to complete its $K-2d$ cell transmissions. The first empty slot left by station $\{1\}$ will reach station $\{2\}$ after the propagation delay which is equal to d slots. This is represented in Figure 2 by a segment of slope -1. Once station $\{2\}$ has completed its transmissions it leaves empty slots to station $\{3\}$. The above process will continue until the *SAT* comes back again to station $\{N\}$ (see epoch t_2 in Figure 2). When this happens, it can be realized that station $\{N\}$ has transmitted $S-K\cdot(N-1)\geq L$ cells and is thus satisfied. Hence, according to the *MAC* protocol, the *SAT* is immediately forwarded to station $\{N-1\}$, $Count(N)$ is set to zero, and process *Y* restarts visiting the same states previously visited in the time interval between t_1 and t_2. Hence, the *SAT* always circulates along the ring without being withheld (i.e., stopped) by any station. This behavior continues over and over again with the result that the Y process is regenerative, and Figure 2 also shows that the *SAT* cycle length is equal to *S* slots.

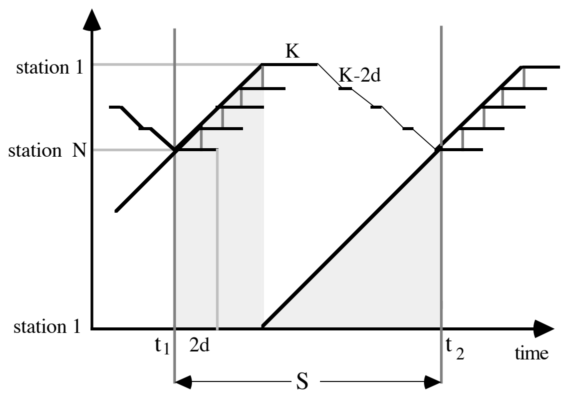

Figure 2: Regenerative cycle when $K\cdot N\geq S,\ K\cdot(N-1)+L\leq S$

Finally, during each *SAT* cycle stations from $\{1\}$ to $\{N-1\}$ transmit K cells, while station $\{N\}$ can access the remaining $S-K\cdot(N-1)\geq L$ empty slots. Since the *SAT* cycle has a duration of *S* slots, the throughput achieved by a generic station $\{i\}$ is given by:

$$\gamma_i = \begin{cases} \dfrac{K}{S} & i=1,2,\ldots,N-1 \\ \dfrac{S-K\cdot(N-1)}{S} & i=N \end{cases} \qquad [3]$$

If, on the other hand, conditions $K\cdot N\geq S,\ K\cdot(N-1)+L>S$ hold, the MetaRing behavior differs depending upon whether $L>2d$ or $L\leq 2d$. However, in both cases, conditions $K \cdot N \geq S,\ K \cdot (N-1) + L > S$ do not always guarantee that any station is satisfied at the *SAT* arrival.

CASE 3: $K \cdot N \geq S,\ K \cdot (N-1) + L > S,\ L > 2d$ - *the SAT stops at station* $\{N\}$ *every round*

Starting from state $\{\mathbf{\Phi}\}$ (see epoch t_1 in Figure 3) and by following the same approach developed for CASE 2, it can be shown that station $\{N\}$ can transmit $2d$ slots before being covered by the busy slots generated by the upstream stations. When this happens, station $\{N\}$ is still not satisfied since $L > 2d$. From then on, the transmissions by station $\{N\}$ will be inhibited until the other stations have transmitted K cells (and are thus satisfied). The *SAT* circulates around the ring without being withheld by the other stations since they all have $Count(j) = K,\ (j = N-1, N-2, \ldots, 1)$. After a complete round around the ring the *SAT* comes back to station $\{N\}$. Due to condition $L > 2d$ the *SAT* is captured by station $\{N\}$ since it still has not performed L cell transmissions. When the *SAT* is withheld by station $\{N\}$, all the others complete their K cell transmissions and only afterwards can station $\{N\}$ transmit the remaining $L-2d$ cells before releasing the *SAT* to station $\{N-1\}$. At this point in time $Count(N) = 0$ and $Count(j) = K$, for $j = N-1, N-2, \ldots, 1$.

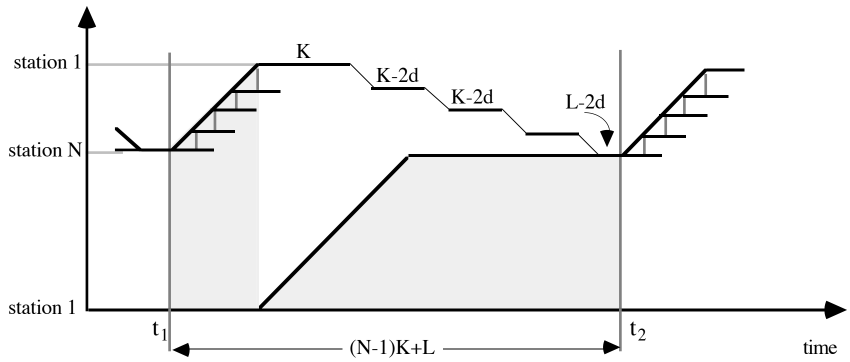

□

Figure 3: Regenerative cycle when $K \cdot N \geq S,\ K \cdot (N-1) + L > S,\ L > 2d$.

Figure 3 shows that at time t_2 a new *SAT* cycle begins during which process Y will go through the same set of states visited in the previous cycle. Specifically, Figure 3 shows that the *SAT* cycle length is $(N-1) \cdot K + L$. Thus

$$\gamma_i = \begin{cases} \dfrac{K}{(N-1)\cdot K + L} & i = 1, .., N-1 \\ \dfrac{L}{(N-1)\cdot K + L} & i = N \end{cases} \qquad [4]$$

CASE 4: $K \cdot N \geq S,\ K \cdot (N-1) + L > S,\ L \leq 2d$ - *the SAT stops at station* $\{N\}$ *every other round.*

Starting from state $\{\mathbf{\Phi}\}$ (see epoch t_1 in Figure 4) and by following the same line of reasoning as in CASE 3, we can readily verify that the *SAT* circulates around the ring without being withheld by the other stations due to $Count(j) = K,\ (j = N-1, N-2, ..., 1)$. After a complete round around the ring the *SAT* comes back to station $\{N\}$.

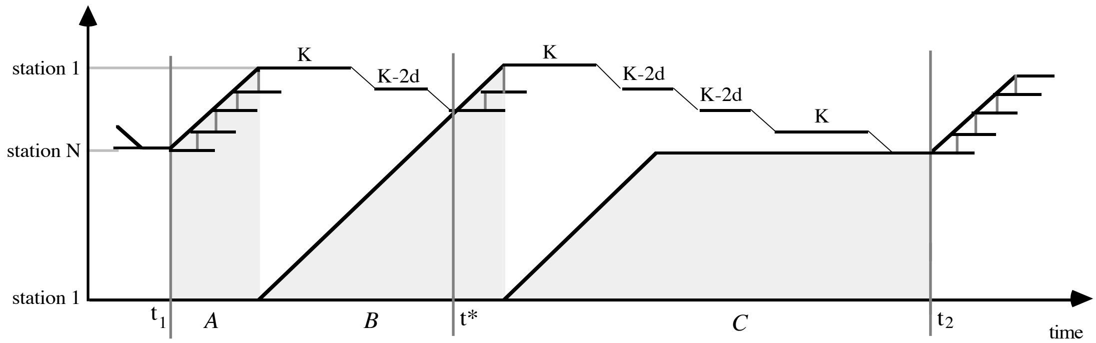

Figure 4: Regenerative cycle when $K \cdot N \geq S,\ K \cdot (N-1) + L > S,\ L \leq 2d$

Due to condition $L \leq 2d$, the *SAT* is not withheld by station $\{N\}$ since it has already been satisfied. Thus, station $\{N\}$ is enabled to transmit again its quota of cells, but it cannot perform any cell transmissions since it is covered by the busy slot string generated by the upstream stations. Clearly, the other stations are all satisfied as well (portion *A* of the *SAT* cycle shown in Figure 4). To derive the structure of the *Y* trajectory during the second round of the *SAT*, let us observe the evolution of the system from the instant the *SAT* reaches station $\{1\}$ in the first round (fraction *B* of the *SAT* cycle in Figure 4). From this point in time onwards, station $\{1\}$ completes *K* consecutive cell transmissions (since it always observes empty slots) and then it gets blocked. After a time interval equal to *d* slots, station $\{2\}$ starts observing empty slots and thus completes the remaining $K-2d$ cell transmissions before getting blocked. This process will continue for the next stations (i.e., stations 3,4,...) until the *SAT*, during the second round around the ring, hits the last station which has already completed or is still completing the remaining $K-2d$ cell transmissions (see Figure 4). This station will be referred to throughout as *intercepted station* and the time at which it is hit will be denoted throughout by t^* (see Figure 4).

Therefore, although during the second rotation (within a *SAT* cycle) of the *SAT*, all the stations were enabled to transmit, cell transmissions restart from the intercepted station and proceed up to station $\{1\}$. When the *SAT* reaches station $\{N\}$ for the second time (within a *SAT* cycle) it withholds the *SAT* until it has transmitted L cells (epoch t_2 in Figure 4) which occurs when all the other stations have transmitted K cells (see fraction C on the *SAT* cycle in Figure 4). From this time onwards the *SAT* cycle restarts again.

To derive the station throughputs γ_i for a generic station $\{i\}$, let us define, for $i=1,2,...,N$, the following time instants (see Figure 5)

I_i : time at which station $\{i\}$ restarts cell transmissions
F_i : time at which station $\{i\}$ stops cell transmissions

counted from the time, which is assumed to be time 0, that the *SAT* leaves station $\{1\}$. Hence

$$I_i = \begin{cases} 0 & i=1 \\ K+(i-1)d+(i-2)(K-2d) & i>1 \end{cases}$$

where, for $i>1$, the first term takes into account K cell transmissions performed by station $\{1\}$, the second term represents the propagation delay experienced by free slots moving towards station $\{i\}$, while the third term is the number of cell transmissions operated by upstream stations $\{i-1, i-2, ..., 1\}$. Furthermore

$$F_i = \begin{cases} K & i=1 \\ I_i + (K-2d) & i>1 \end{cases}$$

Let us now focus on the second *SAT* rotation around the ring. Once the *SAT* has been received by a generic station $\{i\}$ (which is satisfied) it performs two actions. First, $Count(i)$ is cleared and thus cell transmissions by station $\{i\}$ are enabled again; and second, the *SAT* is immediately forwarded to (i.e., it is not withheld by) station $\{i-1\}$. If we denote by SAT_i the time (again measured in slots) at which the *SAT* hits station $\{i\}$ (time 0 is still the time at which the *SAT* leaves station $\{1\}$), we have

$$SAT_i = S - (i-1)d.$$

By indicating with T'_j the number of cell transmissions (this number is equal to the value of $Count(j)$) performed by station $\{j\}$ between t_1 and t^* (see Figure 4) we have for station $\{1\}$

$$T'_1 = \min(K, S)$$

while for stations $i = 2,3,4,...,N$

$$T'_i = \begin{cases} K & if \;\; SAT_i \geq F_i \\ (SAT_i - I_i + 1) + 2d & if \;\; I_i \leq SAT_i < F_i \\ 2d & if \;\; SAT_i < I_i \end{cases}$$

The term "+1" in brackets on the right hand side takes account of the fact that, according to the *MAC* protocol the $Count(\cdot)$ variable is cleared when the *SAT* is forwarded to the next station rather than when the *SAT* is received. Figure 5 also reports the *SAT* position as a function of time (hairline). Figure 5 highlights that we can partition the stations at most into the following three groups: a) stations which were able to transmit K cells (those above the intersection point Q), b) stations which transmitted $2d$ cells only (those below Q), and c) the only station which transmitted a number of cells between $2d$ and $K-2d$ (see station $\{3\}$ in Figure 5). These considerations justify portion B in the cycle in Figure 4. The station with the highest index belonging to group a) or, if available the station belonging to group c) is obviously the *intercepted station* defined above.

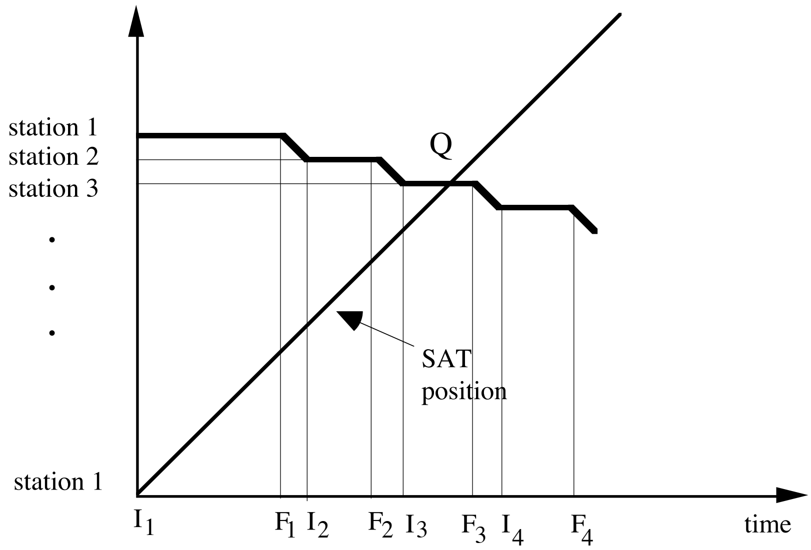

Figure 5: *Y* process evolution after the *SAT* left station $\{1\}$

By denoting with T_i'' the number of cells transmitted by station $\{i\}$ between t^* and t_2 (see Figure 4), we have

$$T_i'' = \begin{cases} K & i = 1, \ldots, N-1 \\ L & i = N \end{cases} \qquad [5]$$

Once station $\{N\}$ has completed *L* cell transmissions, the *SAT* will be forwarded to station $\{N-1\}$ and a new cycle will start again. The throughput achieved by station $\{i\}$, is computed by dividing the number of cells transmitted by station $\{i\}$ in a *SAT* cycle by the cycle length. From Figure 4 we have

$$\gamma_i = \frac{T_i' + T_i''}{\sum_j T_j' + (N-1)K + L}$$

where the summation in the denominator indicates the global number of cells transmitted by all the stations during the first *SAT* rotation while the remaining terms represent the total number of cells transmitted during the second *SAT* rotation. After some algebraic manipulations, the following expression for the throughput achieved by station $\{i\}$ can be obtained

$$\gamma_i = \frac{T_i}{K \cdot A + (B - A - 1)(S + 1 - K \cdot A) + (N - B + 1) \cdot 2d + (N - 1) \cdot K + L}$$

[6]

where:

$$A = \left\lfloor \frac{S}{K} \right\rfloor \text{ and } B = \left\lceil \frac{S-2d+K}{K} \right\rceil$$

$$T_i = \begin{cases} \left.\begin{array}{ll} 2K & i \le A \\ S + K(1-A) + 1 & i = A+1 \ne B \\ 2d + K & i \ge B \end{array}\right\} & for \;\; i = 1,2,\ldots,N-1 \\ 2d + L & for \;\; i = N \end{cases}$$

[7]

From [7] it appears that stations can be grouped into at least four throughput classes. During the first *SAT* rotation there are at least three throughput classes, each associated with a contiguous group of stations characterized by the number of cells transmitted:

> CLASS 1: $\min(K,S)$ cell transmissions for the most upstream (with respect to data transmission) stations, i.e., for the stations verifying the condition $SAT_i \ge F_i$ (station $\{1\}$ always belongs to this class and is certainly the only one when $K > S$).
>
> CLASS 2: $\left(SAT_i - I_i + 1\right) + 2d$ cell transmissions for at most one intermediate station between stations belonging to class 1 and class 3.
>
> CLASS 3: $2d$ cell transmissions for the most downstream (with respect to data transmission) stations.

During the second *SAT* rotation, stations $\{1,2,\ldots,N-1\}$ transmit K cells while station $\{N\}$ transmits L cells. By grouping all the stations into at most three classes during the first *SAT* rotation and into two classes (see [5]) during the second *SAT* rotation, it is easily verified that all the stations can be grouped into at most four classes. Station $\{1\}$ always belongs to the class with the highest throughput, while station $\{N\}$ belongs to the class with least throughput.

By using the above closed formulae we derive the throughput achieved by each station for a MetaRing configuration with $S = 150$, $N = 10$ stations, $d = 7$ and for several values of K and L. For each value of K and L we have to apply the appropriate formula (see Figures 6 - 9).

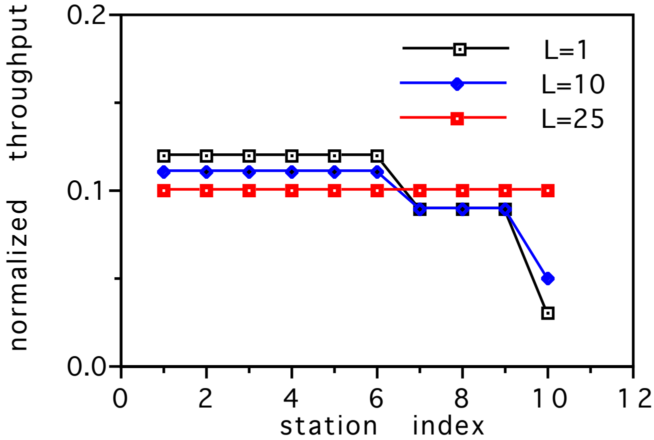

Figure 6: Station throughputs for K=25

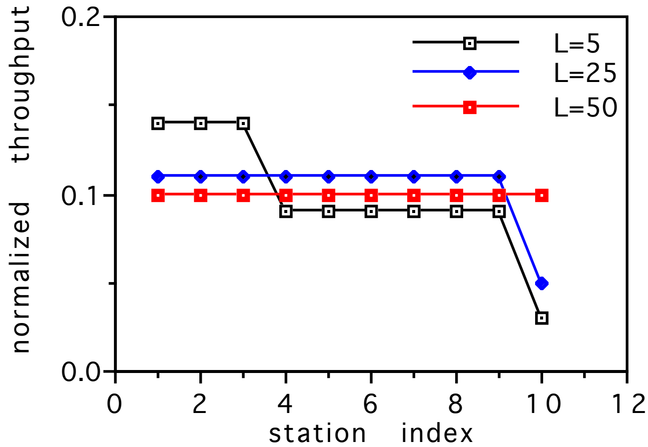

Figure 7: Station throughputs for K=50

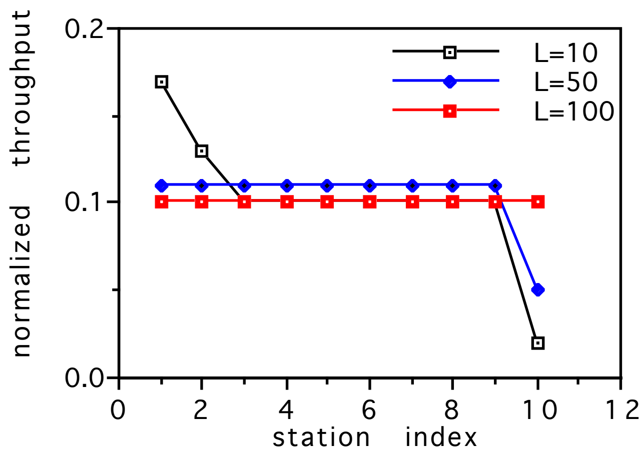

Figure 8: Station throughputs for K=100

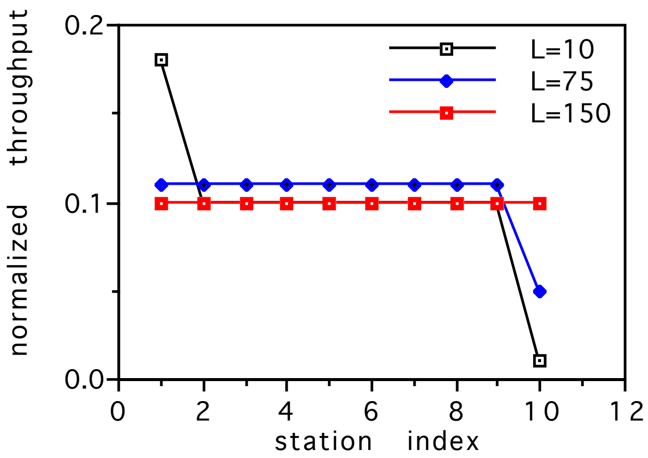

Figure 9: Station throughputs for K=150

These figures clearly show that for $K = L$ all the stations achieve the same throughput. On the other hand, for $K > L$ the least favoured station is the last one, while the most favoured stations are those close to the first station. In particular, note that in Figure 8, the curve for $K = 100$ and $L = 10$ shows four different station throughput values. This corresponds to the fact that for these K and L values all stations must be partitioned into four groups, each with a different value for T_i. As a consequence [6] gives rise to four different values for the station throughputs. On the basis of the previous closed formulae for the aggregate and station throughputs we can draw (at least for the scenario in Figure 1) the following conclusions.

1. Relation [1] shows that for certain values of the K parameter the MetaRing network is underutilized (e.g., $N = 1$, $K < S$). The value of L does not influence the aggregate throughput. Formula [1] also shows that the quantity $(NK)/S$ is an upper bound for the aggregate throughput.
2. Station throughputs, and hence fairness, are very sensitive to the K and L values. Small variations in the parameter values may cause a significant variation in the station throughputs. This occurs when we pass from the range of use of a formula to that of another. For example, with $S = 150$, $N = 2$, $K = 149$, $L = 1$, and $d = 37$ we are in the field of applicability of formula [3] which yields $\gamma_1/\gamma_2 = 149$, while, with $K = 150$, $L = 1$, we are in the range of use of formula [6] which leads to $\gamma_1/\gamma_2 = 4$. In fact, in the former case the *SAT* rotates freely around the ring and the 150 slots (S) are divided by the two stations in the ratio of 149 slots for station $\{1\}$ and the remaining 1 slot for

station $\{2\}$. In the latter case, the *SAT* stops cyclically at station $\{2\}$ and this allows a more balanced ratio between the throughputs of the two stations.

3. By increasing the K/L ratio, the difference between the throughputs achieved by the first and the last station increases. The growth is not linear. Instead, a discontinuity can be observed while passing from the range of use of a formula to the range of use of another. For example, for a network configuration of two stations $(N=2)$, from Table 1 we see that $\gamma_1/\gamma_2 = 3.63$ when $K=300$ and $L=50$ (formula [6] holds), while $\gamma_1/\gamma_2 = 4$ when $K=300$, $L=75$ (formula [4] holds).

L \ K	1	5	10	25	50	75	100	125	150	175	200	300
300												1
200											1	1.5
175										1	1.14	1.71
150									1	1.17	1.33	2
125								1	1.2	1.4	1.6	2.4
100							1	1.25	1.5	1.75	2	3
75						1	1.33	1.67	2	2.33	2.66	4
50					1	1	2	2.02	2.42	2.62	2.83	3.63
25				1	1	1	2	5	3	3.29	3.54	4.56
10			1	1	1	1	2	5	3.57	3.87	4.15	5.33
5		1	1	1	1	1	2	5	3.6	4.1	4.43	5.71
1	1	1	1	1	1	1	2	5	4	4.33	4.67	5.98

Table 1: Ratio γ_1/γ_2 as a function of K and L for a network configuration with $N=2$, $S=150$

4. The equality $K = L$ provides the maximum degree of fairness between stations. Since the aggregate throughput, for fixed K, does not increase when L decreases we can conclude that the choice $K = L$ is the most appropriate for the network scenario considered herein (see Table 1). However, this is not always true. In fact, let us consider the scenario in Figure 10, analyzed via simulation, where stations $\{1\}$ and $\{2\}$ send traffic to the gateway while all the other stations send traffic to the one after. According to this scenario, stations $\{3\}$, $\{4\}$ - $\{12\}$, and $\{1\}$ always observe empty slots. On the other hand, station $\{2\}$ is covered by the traffic sent by station $\{1\}$. However, due

to the SAT algorithm, a station refrains from transmitting if quota K has been reached. From the above observations it follows that:

- the aggregate throughput is less than or equal to the maximum allowed aggregate throughput which, in this scenario, is equal to 11;
- for fixed L, the aggregate throughput increases when K increases (see Table 2). This can be justified by observing that, when K increases the opportunity for a station to stop transmitting (before receiving the SAT) decreases.

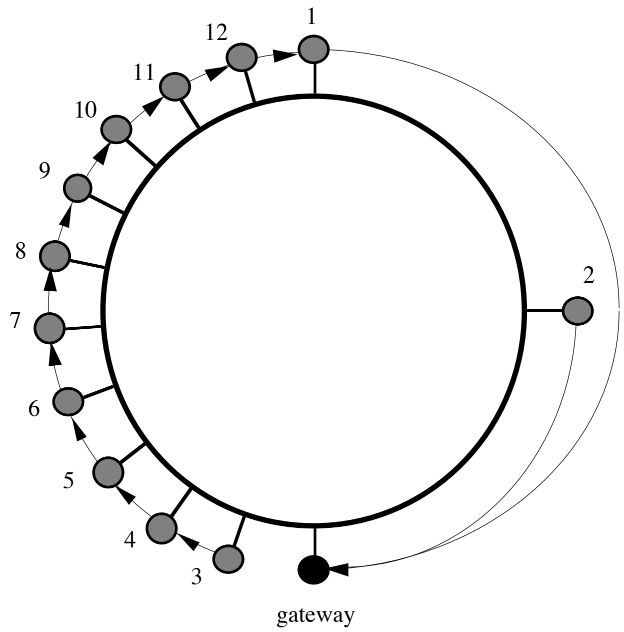

Figure 10: Alternative scenario ($N = 12$, $S = 150$)

L \ K	1	5	10	25	50	75	100	125	150	175	200	300
300												6
200											6	7
175										6	6.33	7.32
150									6	6.38	6.71	7.67
125								6	6.45	6.8	7.15	8.06
100							6	6.56	7	7.36	7.67	8.5
75						6	6.72	7.25	7.67	8	8.27	9
50					4	6	7.67	7.68	8.01	8.23	8.38	8.84
25				2	4	6	7.67	9.33	8.51	8.67	8.79	9.2
10			0.8	2	4	6	7.67	9.33	8.81	8.95	9.01	9.42
5		0.4	0.8	2	4	6	7.67	9.33	8.92	9.04	9.15	9.51
1	0.08	0.4	0.8	2	4	6	7.67	9.33	9	9.12	9.24	9.57

Table 2: Aggregate throughput for the scenario in Figure 10

4. Analysis in Underload Conditions

In this section we analyze the MetaRing network behavior when the user-generated traffic, normalized to the maximum aggregate throughput achievable by the stations in asymptotic conditions, is less than one. Hereafter, the above ratio will be termed *Offered Load* (OL). This analysis, as well as the investigation in overload conditions (i.e., $OL > 1$) which will be presented in the next section, was performed by simulation. To this end, we developed a stochastic queueing model of the MetaRing and solved it via the *RESQ* simulation tool ([GORD92]). The network scenario analyzed is the same as the one considered in the previous section. Details on the simulation model can be found in [MOTT94] while the parameter values which completely define our model are reported in Table 3.

- Stations were spaced equally along half of the ring
- Capacity of each ring = 150 Mbps
- Slot size = 53 octets
- Slot duration *(ts)* = 2.82 μsec
- Length of each ring (S)= 150 slots ≈ 85 km
- Ring latency 423 μsec
- Total number of stations (N) = 10

Table 3: Network parameter values

To have acceptable simulation times we decided, for most of the experiments, to set the channel bandwidth of each ring equal to 150 Mbps, which is in the range of the channel bandwidth for the current MANs. The simulation time of each run was chosen so that the resulting confidence intervals were in the order of 5%[1].

We performed a simulation analysis with several workloads, each characterized, in addition to the offered load, by the distributions of the message interarrival times and message lengths. In our experiments we considered exponential (coefficient of variation C equal to one) hypoexponential $(C < 1)$, and hyperexponential $(C > 1)$ interarrival time distributions [KLEI75]. Furthermore,

[1] The confidence level was chosen to be equal to 90%.

the message length *Msg* was assumed to be constant in each experiment and equal to 1, 4, and 8 cells.

In underload conditions the performance measure of interest in our analysis is the *station average access delay*, i.e. the average time interval between a cell arrival at the station queue and its transmission onto the media. Obviously, in a fair network with equally loaded stations, each station should experience the same average access delay. Any deviation from this behavior has been assumed as being a measure of the *MAC* protocol unfairness in underload conditions.

The station average access delay depends on several parameters that characterize the workload, the number of active stations (N), the *MAC* protocol parameters (*K* and *L*), and the MetaRing network configuration (i.e., ring length, slot size, etc.). It was thus decided to divide the experiments into the five classes shown in Figure 11.

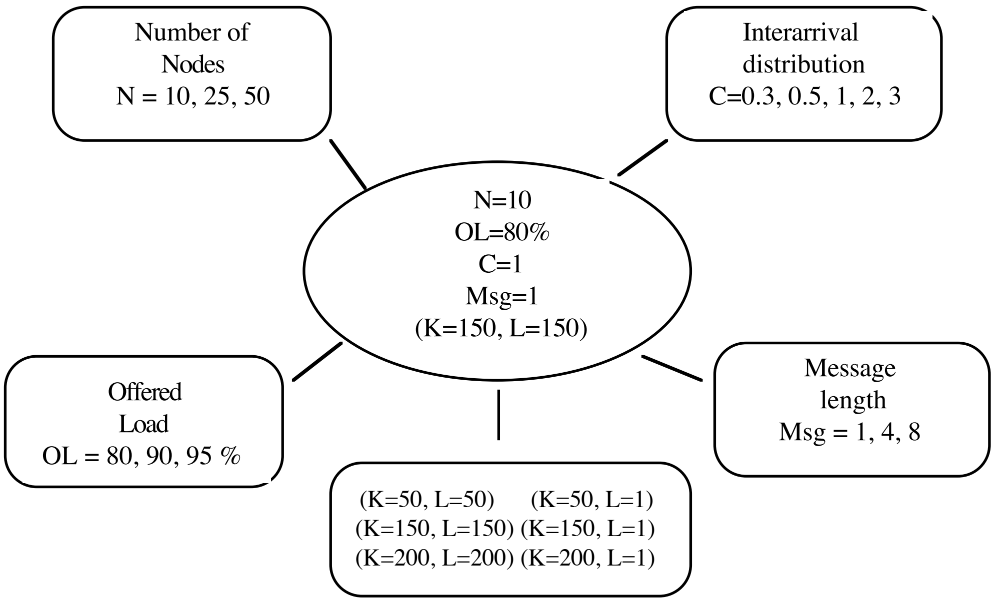

Figure 11: Organization of underload experiments.

Beginning with the following set of parameters $(N = 10, OL = 80\%, Msg = 1, C = 1, K = 150, L = 150)$ all the other experiments were performed by varying the parameters one at a time (except for *K* and *L* which were varied together) between simulation runs. The rest of this section discusses the results.

INFLUENCE OF K AND L

First we considered different couples of values for K and L. For each value of K we considered the cases $L = 1$ and $L = K$ which give rise, in asymptotic conditions, to the maximum unfairness and minimum unfairness, respectively. However, we found that the average access delay experienced by each station does not depend on the K and L values[2]. This behavior can be explained by the fact that the value of K is greater than the average number of cell arrivals at any station during an interval time equal to the average SAT rotation time.

INFLUENCE OF OFFERED LOAD

Figure 12 shows the influence of the offered load (OL) on the average access delay. The initial shape of the curve can be explained by noting that the first group of stations observes a large number of empty slots and thus their station average delays are not significantly affected by the OL considered. On the other hand, the last group of stations (stations 7-10) observes trains of contiguous busy slots with increasing length as the OL increases. This explains the sharp increase in the average access delay in the curve for these stations.

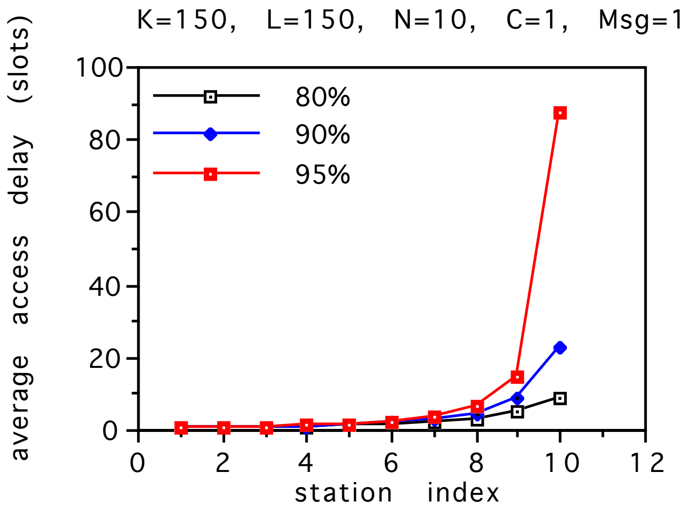

Figure 12: Average access delay vs station index (N=10)

INFLUENCE OF MESSAGE LENGTH

Figure 13, which shows the MetaRing station average access delay as a function of message length, highlights that the MetaRing unfairness increases significantly with message length. This can be intuitively justified by observing that one message arrival contributes to the aggregate OL with a number of cells equal to the message length. Therefore each station will transmit a batch of cells rather

[2] The same results hold for OL=90% and OL=95%

than a cell at a time with the result that the correlation between consecutive busy slots increases. Hence, stations approaching station $\{10\}$ observe trains of busy slots which tend to increase as the station index increases.

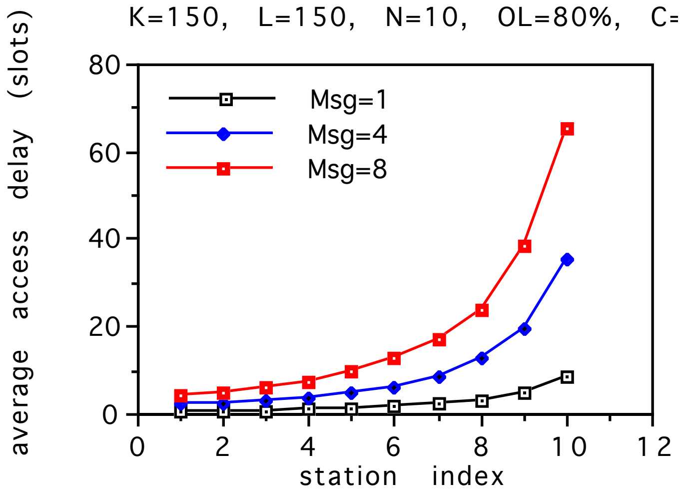

Figure 13 - Influence of Message length.

INFLUENCE OF INTERARRIVAL DISTRIBUTION

Figure 14 reports the MetaRing station average access delay as a function of the interarrival distribution. This figure shows that increasing the C parameter significantly affects the average access delay at least for the offered load underway.

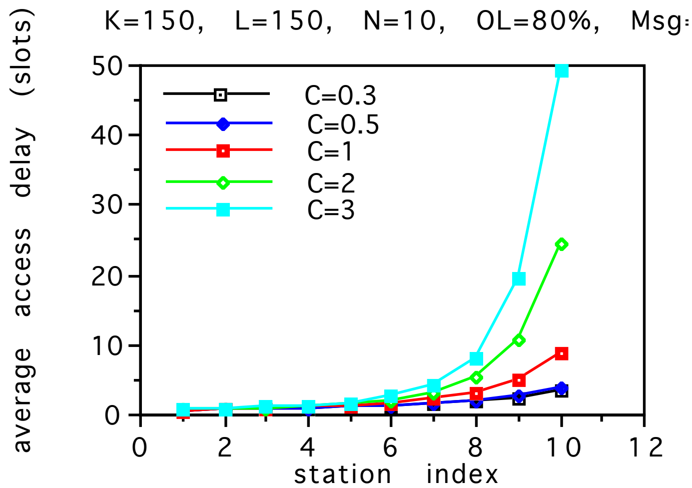

Figure 14 - Influence of Interarrival Distribution.

The shape of the curves can be justified by observing that when C increases, cells tend to arrive in groups. Hence, the considerations already made for $Msg > 1$

hold. On the other hand, when C approaches zero, arrival processes at the various stations tend to become periodic. In fact, the correlation between consecutive slots made busy by the same station decreases, so that the length of busy trains observed by each station decreases. The curve of the average access delay consequently tends to become flat as C tends to zero.

INFLUENCE OF N

To further deepen the analysis of the average access delay we performed a set of experiments in which we varied the number of active stations while leaving the ring length constant. Specifically, we performed two experiments by assuming $N = 25$ and $N = 50$, respectively. Results obtained ([ANAS96]) show that the dependency of the average access delay from the station index is the same as the one already shown in Figure 12 ($N = 10$). However, the difference between the average access delay experencied by the first station and that experienced by the last station increases as N increases.

WORST CASE ANALYSIS

The analysis clearly shows that the average station access delay depends on several parameters. Specifically, it depends on the number of stations, the station index, and the workload parameter values. Hence it was decided to analyze the MetaRing behavior in the worst possible conditions for the values assumed by these parameters.

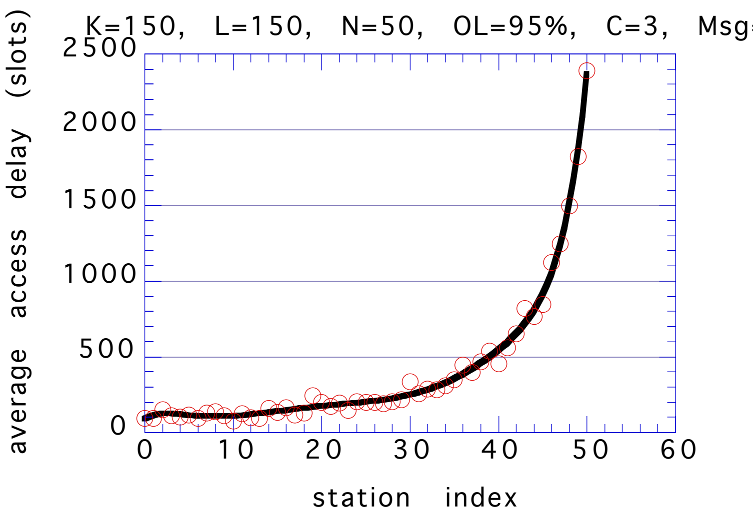

Figure 15: Average access delay vs station index in a worst case set of network and workload parameters.

Figure 15 reports the results of this analysis: circles are the results obtained from simulation, while the curve was obtained by a ninth order polynomial interpolation. The shape of the curve for station indices in the range 1 to 20 is due to the polynomial order used for the interpolation.

In conclusion, we can observe that stations experience varying average access delays, depending on their position on the (semi)ring, even when $K = L$. Thus, at least in the network scenario analyzed, in underload conditions, the MetaRing *MAC* protocol suffers from unfairness.

ANALYSIS AT HIGHER SPEEDS

Since the MetaRing *MAC* protocol was designed to work at speeds in the Gbit/sec region we decided to broaden our analysis to verify whether it can exploit these rates efficiently.

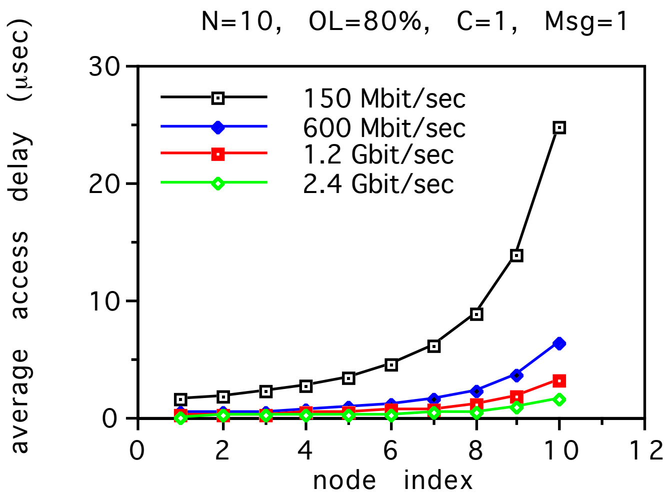

Figure 16: Average access delay vs station index for several channel bandwidths.

Figure 16 plots the average access delay (in µsec) vs station index for several channel bandwidths. To keep the comparison meaningful K and L were varied accordingly. Specifically, we set K and L both equal to the ring length measured in slots. Furthermore, in all the experiments the slot size was assumed to be equal to the *ATM* cell length (i.e., 53 bytes). The shape of the curves can be justified by observing that by increasing the channel bandwidth with *OL* unchanged, the time a cell spends in each station decreases.

5. Analysis in overload conditions

The network scenario is still the same (Figure 1) but the offered load is greater than the maximum achievable aggregate throughput. In overload conditions the performance measure of interest in our analysis is packet loss. In this section we investigate how the packet loss is distributed among stations as a function of the *K* and *L* parameters, *OL*, and the distribution of the workload among stations. The buffer size of stations was chosen in order to have a packet loss less than 1% in underload conditions and also with values of *OL* very close to 100%. By simulation we found that a buffer size of 150 cells can satisfy this requirement. Figure 17 shows the packet loss vs station index for two different *OL* values. The experiments were performed by choosing $K = L = 150$ and a workload uniformly distributed among stations. The packet loss is approximately the same for each station. This can be explained by the fact that, for $K = L$ each station achieves the same throughput.

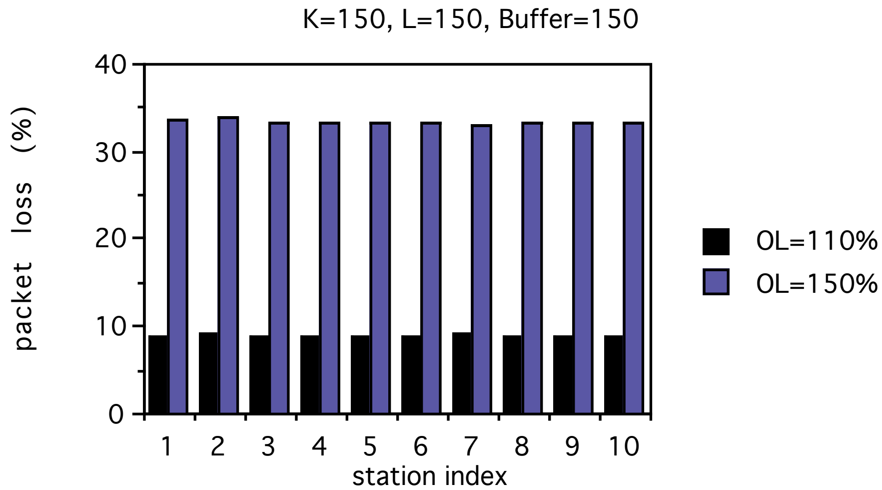

Figure 17: Packet loss vs station position with workload uniformly distributed

In order to further the analysis we studied packet loss as a function of station positions on the ring for *K*=150 and *L*=10. The results are reported in Figure 18 and, for *OL*=150%, they were generally as expected. In fact, for *OL*=150% the ring operates close to saturation and thus the throughput achieved by each station should be approximated by the values reported in Figure 9, for *K*=150, *L*=10. According to this figure, from the throughput standpoint, station $\{1\}$ is the most favoured (i.e., least packet loss) while station $\{N\}$ is the least favoured (i.e., greatest packet loss) while stations $\{3\} - \{9\}$ achieve a constant throughput which is intermediate between those obtained by station $\{1\}$ and station $\{N\}$ (i.e., an intermediate packet loss).

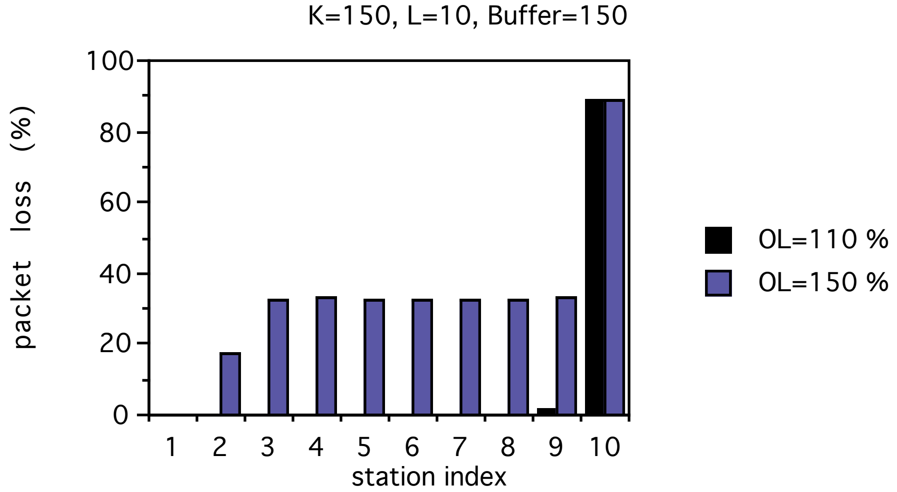

Figure 18: Packet loss vs station position with workload uniformly distributed

The various packet losses experienced by the stations, induced us to study the influence of a group of stations operating in overload conditions (i.e., with an *OL* greater than the maximum achievable throughput) on another group of stations operating in underload conditions. We performed a simulation experiment, with $K = L = 150$, in which the workload was distributed among stations as shown in Figure 19 (which also reports the maximum achievable throughput by each station, i.e. 15 Mbps), and we measured the packet loss of each station. Our results show that the last two stations, although operating in underload conditions, experience a small, but non-zero, packet loss. Clearly with $K = 150$, $L = 10$ we can expect a worse situation as far as the packet loss of stations in underload is concerned. This was, in fact, confirmed by a simulative experiment. We then performed another experiment, with $K = L = 150$, by placing the underload stations upstream of the overload stations. This time no packet loss was recorded in the underload stations. The results could be predicted since in the network scenario analyzed, stations which follow the gateway (i.e., $\{1,2,...\}$) have some advantage over stations which precede the gateway (i.e., $\{N, N-1,...\}$). The fact that the packet loss measured for stations in overload is less than the percentage of workload which exceeds the maximum achievable throughput (i.e., 15 Mbps) by these stations is because they can exploit the portion of the network capacity left free by the stations operating in underload conditions. In any case, the aggregate throughput is never greater than 1, which is the value obtained in asymptotic conditions.

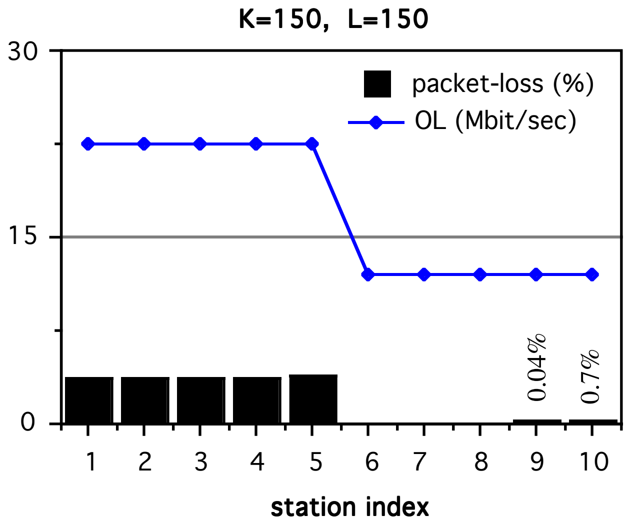

Figure 19: Packet loss vs station position with an unbalanced workload

6. Conclusions

In this paper we have reported some results of an in depth analysis of the asynchronous capabilities of the MetaRing *MAC* protocol by using specific network configurations. From our analysis the following conclusions can be drawn.

ASYMPTOTIC CONDITIONS

- For certain values of the K parameter the MetaRing network is underutilized (i.e., $\rho < 1$).
- Station throughputs, and hence fairness, are very sensitive to K and L values. Small variations in their values may cause a significant variation in station throughputs.
- The maximum degree of fairness is achieved when $K = L$. Since the aggregate throughput, for fixed K, does not increase when L decreases, the choice $K = L$ is the most convenient for the network scenario reported in Figure 1. However, this is not true in general. In fact, for the scenario in Figure 10, for fixed K, the aggregate throughput increases when L decreases (see Table 2).

UNDERLOAD CONDITIONS

- $K = L$ does not guarantee fairness. In fact, when $K = L$ we can observe that stations experience different average access delays, depending on their position along the network.

- Furthermore, the fairness depends on several other factors: offered load, message length, interarrival time distribution, and number of active stations.

OVERLOAD CONDITIONS

- When the workload is uniformly distributed among stations and $K = L$, the packet loss is the same for every station. On the other hand, if $K > L$ the stations closer to the gateway exhibit a packet loss which is greater than the others.
- When the workload is not uniformly distributed among stations (Figure 19), the underloaded stations experience a non zero packet loss.

Finally, the MetaRing performance and fairness are very sensitive to minor changes in parameter settings. However the *SAT* is very simple to implement. So we have a typical trade-off between performance/fairness and simplicity.

This analysis has been further progressed in [PEVA], where different traffic patterns (though with the same network configuration) have been analysed. An analysis regarding alternative station arrangements is under study.

References

[ANAS95] G. Anastasi, L. Lenzini, P. Motta, "Performance Evaluations of a MetaRing MAC Protocol Integrating Video and Data Traffics in an Interconnected Environment", Proceedings INFOCOM '95, Boston (USA), April 1995, pp. 1273-1281.

[ANAS96] G. Anastasi, L. Lenzini, "The MetaRing MAC Protocol: a Performance Analysis", Internal Report 1996.

[CIDO93] I. Cidon, Y. Ofek, "MetaRing, a Full Duplex Ring with Fairness and Spatial Reuse", IEEE Transaction on Communications, Vol. 41, No. 1, January 1993.

[GORD92] R. F. Gordon, P. G. Lowener, E. A. MacNair "The Research Queueing Package Version 3. Language Reference Manual", *IBM Research Report*, Yorktown 1992.

[KLEI75] L. Kleinrock, "Queuing Systems", Volume I: Theory, John Wiley & Sons, 1975.

[MOTT94] P. Motta, "Integrazione dei Servizi nel MetaRing: Valutazione delle Prestazioni", Laurea Thesis (in Italian), University of Pisa, Engineering Faculty, Pisa (Italy) 1994.

[PEVA] G. Anastasi, L. Lenzini, B. Meini, “Performance Evaluation of a Worst Case Model of the MetaRing MAC Protocol with Global Fairness”, to appear on Performance Evaluation.

[OFEK94] Y. Ofek, “Overview of the MetaRing Architecture”, Computer Networks and ISDN Systems, Vol. 26, Nos. 6-8, March 1994, pp. 817-830.

[WU92] H. T. Wu, Y. Ofek, K. Sorhaby, "Integration of Synchronous and Asynchronous Traffic on the MetaRing Architecture and its Analysis", Proceedings of ICC '92.