PROPOSALS FOR IMPROVING FASNET PERFORMANCES

Internal Report C88-50

ottobre 1988

E.Gregori
L.Lenzini

PROPOSALS FOR IMPROVING THE FASNET PERFORMANCES

E. Gregori, L. Lenzini
CNR-Istituto CNUCE
Via S. Maria, 36
56100 Pisa

1.0 **Introduction**

FASNET operates as a synchronous network [LIMB82] using a transmission medium made up of a pair of unidirectional lines (Figure 1). In the current implementation they are of optical fiber. The protocol operates on each line independently. The stations at the beginning of each line (referred to as Head Stations HS) transmit a synchronization code at regular intervals, thus creating a slotted stream, which other stations synchronize on. A slot has an internal structure made up of an Access Control Field (ACF) and an Information Field (IF). IF will be filled with packets by the transmitting stations. In the following, the set of the ACF and IF will be referred to as a frame. When a station puts a packet on a slot, the station itself marks a bit in the ACF called BUSY bit to indicate that the frame is taken.

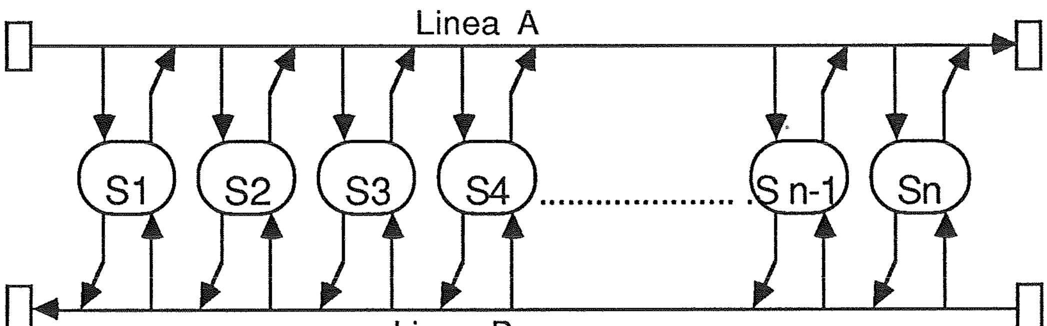

figure 1 - Fasnet network structure

Stations are ordered according to their physical position with respect to the head station and univocally identified by an integer number denoted by i. On line A station 1 is the leftmost and station N is the rigthmost.

In addition to synchronization codes, the head/end stations send START/END bits as described below. The HS related to line A (i.e. S_1)

initiates a cycle (frame with START=1) , then each active station (i) is allowed to use up to P_{max}(i) consecutive[1] slots in this cycle. This prevents stations closest to the HS from unfairly seizing all frames.

When the last station on line A (i.e. S_N) detects an empty frame, it sends a signal back to the first station by marking the END bit in the ACF of the first frame on line B. Upon receiving the END=1, the HS for line A marks the start-of-cycle by means of the START bit in the next frame, thus allowing stations to transmit again. For details on the FASNET access protocol the reader is encouraged to read [LIMB82].

2.0 **Definitions and Claims**

In this section definitions and claims extensively used in the paper are provided.

v = signal propagation speed (m/s)
W = line capacity (bit/s)
L = line length (m)
F = frame size (bits)
M = number of busy stations with downstream traffic,
N = total number of stations
$t_{I,1}$ = transit delay related to the propagation of the optical signal from station 1 to station I
$\Delta_{I,1}$=maximum number of slots that can be generated by the HS in the period $(t-t_{I,1}, t+t_{I,1})$.

Definition A cycle (T_{cl}) is the time between two consecutive START bits.

For the basic FASNET, T_{cl} can be regarded as the sum of two types of delay:

$T_{cl}= T_{bf}+ S_j$ [1] where:

- T_{bf} (bf standing for busy frames) is the time required by the active stations (i.e., stations that have packets ready for transmission) to

[1] In the current implementation of FASNET a station can also operate in another way. A station doesn't have to transmit all its packet consecutively provided that its P_{max} per cycle is not exceeded.

transmit their packets. For T_{bf} the following condition holds:

$0 \leq T_{bf} \leq P_{max}(i)$ [2]

T_{bf} is generally variable since an active station may not need to transmit its full allowance, and not all stations may be active during a given cycle.

- S_j is the sojourn time; that is, the time it takes the HS to discover the end of the cycle. The maximum and the average values of S_j are respectively 2*(L/v)+2*(F/W) and 2*(L/v)+(F/W). In the worst case, Sj is then equal to twice the end-to-end propagation delay, plus twice the slot time (one slot time on average) as each End Station (ES) has to wait until the start of the next slot to set the START or END bits.

Definition The FASNET utilization (U) of the media is defined by the following relation:

$U=T_{bf}/T_{cl}$ [3]

Comment Although there are no collisions, U is not 100 percent because of the sojourn time. This wasted time is becoming a significant portion of the cycle time for MANs with a length in the order of 100 Km. For example, if v=2.5x10**8 m/sec, W=1.2x10**9 bit/sec, L=100 Km, F=1024 bits, P_{max}=16 (constant), N=M=50 then T_{cl}=1.483 msec, U=0.46 (i.e., 46%!!).

Definition $T_{hr}(i)$, the throughput within a cycle of station i, is defined by the following relation:

$T_{hr}(i)=P(i)*F_t/T_{cl}$ [4]

where P(i) [P(i) being less than or equal to $P_{max}(i)$] is the actual number of packets transmitted by station i within the cycle underway, F_t is the length of the slot in seconds (for the sake of precision, F_t should be the IF part of the frame expressed in seconds) and T_{cl} is the length of the cycle underway in seconds.

Claim $T_{hr}(i)$ has a maximum and a guaranteed minimum .
Proof From [4] it follows that $T_{hr}(i)$ is maximum when station i transmits $P_{max}(i)$ packets within a cycle in which all the other stations are silent $T_{hrmax}(i)$. In this case $T_{hr}(i)=P_{max}(i)*F_t/(P_{max}(i)*F_t+S_j)$. Thr(i) is equal to the minimum guaranteed throughput,Thrmin(i), when all the stations, including station i, transmit the maximum number of packets they are allowed to (i.e., $P_{max}(i)$, for each i less than or equal to N and greater than or equal to one).

2.0 **Proposals for improving the FASNET performances**

As can be seen from [1], [2] and [3] U increases (i) as S_j decreases and (ii) as T_{bf} increases.

As far as (ii) is concerned, it is obvious that by controlling P_{max}, stations may influence the value of T_{cl} [LIMB82]. However, this approach is clearly limited. In fact, if a station generates packets at a rate $<P/T_{cl}$, any increase of Pmax, from P onward, will not change the cycle length since packets will be transmitted before a queue can form.

In the following some techniques are proposed for improving utilization by acting on the sojourn time.

For each technique the following figures will be defined to provide some major highlights of performances and effectiveness and to ease the comparison among them.

1-Minimal cycle length

Length of the cycle when there is only one station active on the line with only one packet queued for transmission. This parameter gives an idea about the overhead introduced by the FASNET access scheme under very light load conditions.

2-*Thrmax(i)* and *Thrmax(i)/Thrmin(i)* (see 2.0 definitions)

To provide indications of the capacity of the network in order to dynamically adapt to different traffic configurations.

3-*Conditions for U very close to 1*

To point out the conditions that had to be satisfied in order to obtain effective improvement compared to the basic FASNET

4-*Privileged Stations*

To indicate the introduction of privileges. The complete absence of privileges is obtained when the bandwidth available to a station only depends on its Pmax and on the cycle length.

2.1 Techniques for the Sojourn Time Utilization

The basic idea behind these techniques is to make use of the empty frames that travel in the network during the S_j time in order to transmit packets (if any) which are still queued for transmission.

2.1.2 Nearest seizing of empty slots

This technique is a variation of the one described in [LIMB82]. It is based on the fact that ES(i-1,i); i.e., the number of empty slots passing through station i in the time required by END=1 to travel from station i to i-1, might be a non negligible percentage of $P_{max}(i)$ and as such it must be carefully managed.

To prove that, let us make reference to the parameters used in the example shown in the previous section. If the 50 stations are uniformly distributed along the lines, the distance between adjacent stations is 100/50=2Km. The number of empty slots ES(i-1,i) (for each i between 1 and N) accomodated in 2 Km is equal to:

2*{L(i-1,i)/V}/{F/W} =
2*{(2x10**3)/(2.5x10**8)}/{1024/(1.2x10**9)}=~18

This technique makes use of these extra slots as follows: a generic i station in the WAIT state that observes END=1 starts to fill the empty frames passing by. It is easy to demonstrate that at least ES(i-1,i) slots on the opposite line must be empty. Station i has then allowed Pmax(i)+ES(i-1,i) frames per cycle. If station i needs less than ES(i-1,i) slots, the empty slots left will be used by the next (>i) active station (i.e., a station with frames waiting to be transmitted in the cycle underway). Empty slots are thus sized by the active station(s) nearest (>) to the station which does not use, or partially uses, the quantity of empty slots permanently allocated to it. This technique works efficiently when the traffic load of a station tends to increase as its distance from the head station increases. Empty slots unused by the far end stations are very luckly lost. This technique does not require any additional complexity to the already existing basic FASNET access mechanism. Each station only needs to know the quantity of empty slots it can use in addition to P_{max}.

Minimal Cycle Length : $T_{clmin}=S_j$

Thrmax(i) : $P_{max}(i)/[P_{max}(i)+S_j]$

Thrmax(i)/Thrmin(i): $\Sigma_{i=1}^{N}$ Pmax(i)+ S_j/(Pmax(i)+S_j)

Conditions for U very close to 1

This condition can be obtained when each station (i) needs to transmit (within a cycle) a number of packets equal to $P_{max}(i)$+ES(i-1,i). $P_{max}(i)$ are transmitted in the empty frames which follow START=1. ES(i-1,i) are transmitted on detecting END=1. However, it is not necessary that the extra packets (i.e., the packets beyond Pmax(i)) transmitted during the sojourn time be equally distributed among the N stations. The only condition required is that any empty frame passing over station i must be used by a station downstream from station i itself.

Privileged Stations

This technique tends to privilege downward stations closest to the one which does not use, or partially uses, the extra slots ES(i-1,i) it would be allowed to. Note that the capability of the network to absorb traffic bursts on a station depends not only on the global network load but also on the station's positon. For example, in an unloaded network station 1 has only Pmax(1) slots per cycle whereas station N-1 has Pmax(N-1)+ES(1,N-1).

2.1.1 Farthest seizing of empty slots

Compared to the basic FASNET frame format, this technique requires an extra bit (REQ) in the access field of a frame and some more computation in each station for managing a counter denoted by C(i) (i being the station underway). According to this technique, empty slots are seized starting from the nearest active station to the end station.

To describe the operation of this technique the following times are defined:

- $T_c(i)$: the time at which station i detects the first empty slot after the START bit is detected;

- $T_e(i)$: the time at which station i detects, on the return line, the END bit.

At time $T_c(i)$ the difference

D(i)=E(i)-ES(i-1,i), where E(i)=p(i)-MIN[p(i),Pmax(i)]

is computed by station i. E(i) is zero if station i has a number of p(i) packets queued for transmission equal to or less than Pmax(i). D(i) is greater than zero if station i has to transmit extra packets in addition to Pmax(i)+ES(i-1,i). D(i) is zero or less than zero if station i uses respectively all or a subset of the empty slots given by the ES(i-1,i) function. Packets arriving at station i after $T_c(i)$ will be managed in the next cycle. Station(i) will notify its extraload to the upward stations setting the RQ=1 in the first D(i) frames on line B. The condition D(i) $<P_{max}(i)-1$ limits the extra load to Pmax(i)-1 but will avoid any collision between stations in setting the RQ bit. Henceforth, if the previously defined condition is applied, station i may transmit in a cycle up to 2*Pmax(i)+S(i-1,i) packets. Other ways to avoid congestion on the RQ bit may be easily defined.

Once determined D(i), staion i takes the following actions:

C(i)=0;
if D(i)>0 **then** flag XLOAD is set ;
else [D(i) less than or equal to zero] flag NOXLOAD is set;

After station i has transmitted up to Pmax(i) packets it monitors the return line for keeping track of the REQ bits issued by the higher numbered stations (i.e., the ones downstream from i). REQ=1 in the access field of a frame received by station i indicates that a downstream extraload packet is queued for transmission on line A.

On detecting REQ=1 during the interval time $[T_c(i),T_e(i)]$, station i takes the following actions:

if [D(i)>0 & XLOAD] **then** C(i)=C(i)+1; REQ bit unchanged;

if [D(i)<0 & NOXLOAD] **then** D(i)=D(i)+1; REQ bit changed;

if [D(i)≥0 & NOXLOAD] **then** no action is taken on: C(i), D(i) and REQ bit;

At time Te(i) station(i) starts transmitting up to ES(i-1,i) packets. Now C(i), if>0, represents the total number of packets which have not been transmitted yet in the cycle underway at the stations which are downstream from i.

ES(i-1,i) frames after Te(i) station(i) start executing the following procedure upon detecting an empty slot:

if (XLOAD) **then**
{**if** [C(i)=0] **then** use the empty slot;
if [C(i)>0] **then** C(i)=C(i)-1; do not use the empty slot;
}
else (do nothing)

The above procedure boils down to this: station i decrements the C(i) counter by one for each empty slot that passes on line A, since this empty slot will go to serve one of the downstream queued packets waiting to be transmitted. This empty slot will be captured by the first downstream station which has XLOAD&C(i)=0. If the flag NOXLOAD is set, no action is performed. Thus the C(i) count, if greater than zero, is at all times equal to the exact number of packets that are queued for access in the stations downstream from it.

Minimum Cycle Length
(the same as the previous technique)

Thrmax(i)
(the same as the previous technique)

Thrmax(i)/Thrmin(i):)
(the same as the previous technique)

Conditions for U very close to 1
(the same as the previous technique)

Privileged Stations

The privileged stations are now those close to the end station.

2.2 Techniques for the sojourn time minimization

In this section a technique for drastically reducing the sojourn time are described. This technique acts on the length of the cycle. In some significant cases, cycles with an UF very close to 100% are obtained.

2.2.1 End Cycle set by intermediate nodes

The idea behind this technique is to anticipate the END=1. If a station(i) knew the exact amount of frames required in this cycle by the following stations, it could know if the HS generates enough frames to satisfy the requirements of stations(i+1,N-1) during the round trip delay 1,i. Therefore, provided that the requirements of each station are known , or at least estimated, the END=1 can be set by intermediate stations without altering the network fairness.

Definitions:

$S_{max}(I)=\Sigma_{i=I+1}^{N-1} Preq(i)$ = Maximum number of slots required by Stations I+1, N-1 in this cycle.

$\underline{I}$ =the minimum I such that $\Delta_{I,1} \leq S_{max}(\underline{I})$.

It is easy to demonstrate that at least one station I exists such that $\Delta_{I,1} \leq S_{max}(I)$, $1 \leq I \leq N$.

Claim $\underline{I}$ can set the END=1 when the first empty slot goes towards the station $\underline{I}$+1without altering the network fairness.

Demonstration If t is the instant in which the first empty slot overtakes $\underline{I}$ then, because of Fasnet characteristics, stations from 1 to $\underline{I}$ included can't use this cycle's slots anymore. Furthermore, the slots generated by the HS in the period $(t-t_{I,1}, t+t_{I,1})$ match this cycle requirements of stations $\underline{I}$

+1, ..,N-1.

Note that :

$t-t_{I,1}$ is the time when the HS has generated the slot that first overtakes the still empty station **I**.

$t+t_{I,1}$ is the HS reception time of the END bit.

It is important to point out that this method is completely insensitive to traffic variations in stations 1,**I** . Another nice characteristic is that the minimum cycle time is $2*t_{I,1}$ and if the network is lightly loaded **I** can be very close to 1.

How to Estimate I

Several approaches can be devised to estimate **I**. Of course, the optimum will be a knowledge of the station requirements on a per cycle base. This is not both easy to do and very effective. In general, every FASNET station will have enough bandwidth available to perform a good statistical multiplexing; furthermore, the cycle period is some milliseconds. From the previous points it comes that there is no need for update of requirements on a per cycle base. Policies where station requirements are updated only when the difference with the Preq previously notified exceed a certain threshold can be applied. In the simplest method the HS indicates in the Start cycle frame the **I** . Every station will notify significant variations to of its requirement to the HS so that the HS can accurately estimate the **I**. This information (station requirement) is communicated to the HS by means of higher level protocols. Some other guards, such as the number of free slots at the end of a cycle (it can be measured by station N), can trigger the Preq exchange algorithm. Other high level algorithms where **I** is computed in a distributed way can be designed. As a final point, it is important to remember that we can tolerate overestimations (some unused slots in the cycle) whatever algorithm has been chosen but we must avoid underestimation to maintain the network fairness.

Minimal cycle length $T_{clmin} = 2*t_{2,1}$

The only lower limit to the cycle length is that at least two station must be involved (END=1 set by station 2) in order tro have a cycle.

Thrmax(i) $Pmax(i)/ \Delta_{I,1}$

Thrmax(i)/Thrmin(i)

$\Delta_{\max \underline{I},1} / \Delta_{\underline{I},1}$ where max $\underline{I}$ is the estimated I when $\forall$ i $\geq \underline{I}$ Preq(i)=Pmax(i).

Conditions for U very close to 1
Every time $S_{max}(\underline{I})$ has been correctly estimated the U is close to 1. The only source of error is caused by the discreteness of the $\Delta_{\underline{I},1}$.This error is negligible if $t_{\underline{I},\underline{I}-1} \ll T_{cl}$.

Privileged stations There are no privileged stations.

2.2.2 END=1 set by the End Station beforehand

Compared to the basic FASNET protocol, this approach requires:

- 1 more control bits, denoted by REQ, in the Access Control Field of each frame;
- one counter, denoted by C(N), which is managed by the end station (N).

The protocol works as follows: when station i receives a frame for transmission on line A, it sets the REQ bit to one in the Access Control Field of the first slot passing above i (i.e. line A). This bit indicates to the end station that an upstream station (from N) has a frame ready to be transmitted on line A. When station i receives the START bit of the cycle underway, it can transmit in this cycle only the request queued up until this moment. Any new request arriving at station i after the START bit is received should be transmitted in the next cycle. Therefore, this method introduces an additional delay which is always less then the length of the cycle underway.

Any time station N receives REQ=1, it increments C(N) by one. When station N receives the START bit, C(N) contains the total number of frames that, in the current cycle, follows the START bit itself. When this event occurs, station N computes the difference D=C(N)-2*ES(N,1) and sets to zero the counter C(N). If D is greater than, or equal to zero, the sojourn time of the current cycle can be eliminated by simply demanding the end station to issue the END bit after D slots from the START bit. If D is less than zero, station N transmits the END bit according the basic FASNET protocol.

Minimal cycle length $T_{clmin} = S_j$

Thrmax(i) Pmax(i)/ T_{clmin}

Thrmax(i)/Thrmin(i) T_{clmax} /T_{clmin}

Conditions for U very close to 1 Every time that $T_{cl} > S_j$ U is 1

Privileged stations There are no privileged stations.

REFERENCES

[LIMB82] J.O.Limb, C.Flores, "Description of Fasnet - A Unidirectional Local Area Network Communications Network," BSTJ V 61, 7, p 1413, Sept 1982.