

Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCS-88-14

1988-04-01

Performance Models for NOAHNET

Gurudatta M. Parulkar, Adarshpal S. Sethi, and David J. Farber

Noahnet is an experimental flood local area network with features such as high reliability and high performance. Noahnet uses a randomly connected graph topology with four to five interconnections per node and a flooding protocol to route messages. In Noahnet flooding, the routing of a message from a source to the destination node is a two step process: flooding-growth and flooding-contraction. During the growth of flooding, the message propagates to every node which is not occupied with a message and is reachable from the source node. During the contraction of flooding, the nodes that became occupied during the growth... [Read complete abstract on page 2.](#)

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research



Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Parulkar, Gurudatta M.; Sethi, Adarshpal S.; and Farber, David J., "Performance Models for NOAHNET" Report Number: WUCS-88-14 (1988). *All Computer Science and Engineering Research*. https://openscholarship.wustl.edu/cse_research/771

Department of Computer Science & Engineering - Washington University in St. Louis
Campus Box 1045 - St. Louis, MO - 63130 - ph: (314) 935-6160.

Performance Models for NOAHNET

Gurudatta M. Parulkar, Adarshpal S. Sethi, and David J. Farber

Complete Abstract:

Noahnet is an experimental flood local area network with features such as high reliability and high performance. Noahnet uses a randomly connected graph topology with four to five interconnections per node and a flooding protocol to route messages. In Noahnet flooding, the routing of a message from a source to the destination node is a two step process: flooding-growth and flooding-contraction. During the growth of flooding, the message propagates to every node which is not occupied with a message and is reachable from the source node. During the contraction of flooding, the nodes that became occupied during the growth of flooding become unoccupied again. Nodes on unsuccessful paths become unoccupied in a relatively short time compared to the nodes on the successful path. The purpose of this paper is to present two analytical performance models which we have designed to understand the load-throughout behavior of Noahnet. Both models assume slotted Noahnet operation and also assume the if k messages attempt transmission in a slot, the network gets divided into k partitions of arbitrary sizes - one partition for each message. First, we show that the average number of successful messages in a slot given k attempted transmissions is $(M-k)/(N-1)$, where N is the number of nodes in the network and M is the number of nodes out of N that participate in the flooding of k messages. This is an interesting result and is used in both models to derive the load-throughout equations. Each model is then presented using a set of assumptions, derivations of load-throughout equations, a set of plots, and the discussion of results. Models one and two differ in the way they account from retransmissions. Model two helps study the effect the retransmission probability on the performance of the network. The results from these models suggest that the maximum throughput of Noahnet is always less than one message per slot. Also the network is unstable in the sense that the throughput increases with the load only up to a certain threshold value of load; beyond that the throughput starts decreasing with the load. Model two suggests that Noahnet is essentially a contention system, and to get the maximum throughput, the load should be such as to given the optimal contention.

PERFORMANCE MODELS FOR NOAHNET

**Gurudatta M. Parulkar, Adarshpal S. Sethi
and David J. Farber**

WUCS-88-14

April 1988

**Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899**

PERFORMANCE MODELS FOR NOAHNET

Gurudatta M. Parulkar*

Department of Computer Science
Washington University, St. Louis, MO 63130
guru@flora.wustl.edu

Adarshpal S. Sethi

Department of Computer and Information Sciences
University of Delaware, Newark, DE 19711
sethi@dewey.udel.edu

David J. Farber

Director, Distributed Systems Laboratory
Department of Computer Science and Electrical Engineering
University of Pennsylvania, Philadelphia, PA 19104-6389
farber@cis.upenn.edu

*This work was done at the University of Delaware as part of this author's Ph.D. thesis

Abstract

Noahnet is an experimental flood local area network with features such as high reliability and high performance. Noahnet uses a randomly connected graph topology with four to five interconnections per node and a flooding protocol to route messages.

In Noahnet flooding, the routing of a message from a source to the destination node is a two step process : flooding-growth and flooding-contraction. During the growth of flooding, the message propagates to every node which is not occupied with a message and is reachable from the source node. During the contraction of flooding, the nodes that became occupied during the growth of flooding become unoccupied again. Nodes on unsuccessful paths become unoccupied in a relatively short time compared to the nodes on the successful path.

The purpose of this paper is to present two analytical performance models which we have designed to understand the load-throughput behavior of Noahnet. Both models assume slotted Noahnet operation and also assume that if k messages attempt transmission in a slot, the network gets divided into k partitions of arbitrary sizes - one partition for each message.

First, we show that the average number of successful messages in a slot given k attempted transmissions is $(M-k)/(N-1)$, where N is the number of nodes in the network and M is the number of nodes out of N that participate in the flooding of k messages. This is an interesting result and is used in both models to derive the load-throughput equations.

Each model is then presented using a set of assumptions, derivations of load-throughput equations, a set of plots, and the discussion of results. Models one and two differ in the way they account for retransmissions. Model two helps study the effect of retransmission probability on the performance of the network.

The results from these models suggest that the maximum throughput of Noahnet is always less than one message per slot. Also the network is unstable in the sense that the throughput increases with the load only up to a certain threshold value of load; beyond that the throughput starts decreasing with the load. Model two suggests that Noahnet is essentially a contention system, and to get the maximum throughput, the load should be such as to give the *optimal* contention.

1 INTRODUCTION

Noahnet is an experimental flood local area network with features such as high reliability and high performance. Reference [1] is an introduction to the Noahnet architecture, its implementation of flooding, and also its message format. Reference [2] considers some of the other design issues of Noahnet, such as flood control, functions and design of a node, and expected performance of Noahnet. Reference [3] includes the most up-to-date and complete state diagram description of the Noahnet flooding protocol. The purpose of this paper is to present two performance models which we have designed to predict the load-throughput behavior of Noahnet. Model one is a simple load-throughput model with no retransmissions. Model two is also a load-throughput model but accounts for retransmissions, and it is more general. In fact, model one is a special case of model two and is included in this paper for its simplicity.

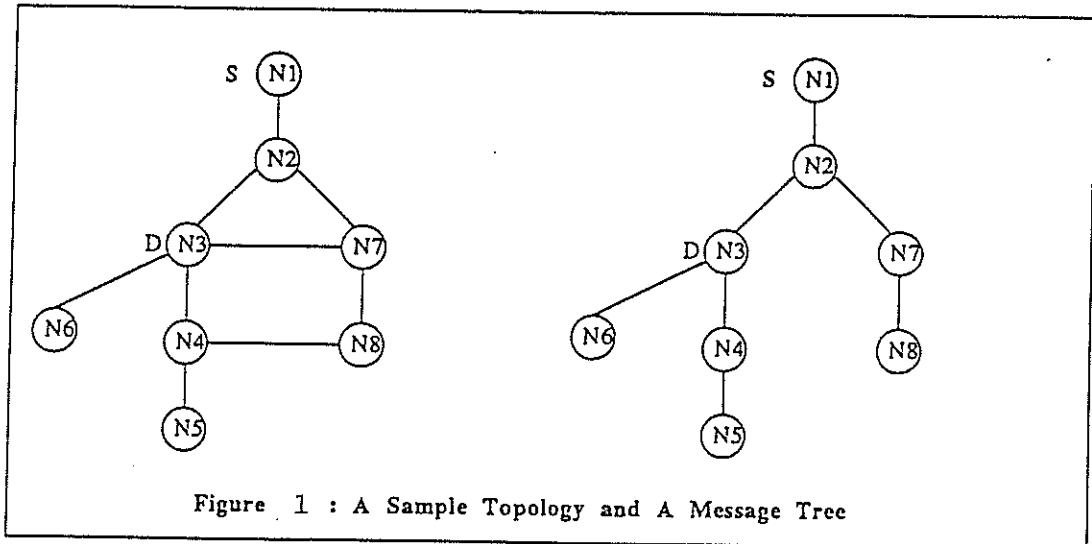
These models suggest that Noahnet's load-throughput behavior is similar to that of a contention system such as CSMA and CSMA/CD networks. In a contention system, throughput increases with load only up to a certain threshold value, beyond which the throughput starts decreasing with increasing load. The threshold value of load corresponds to an optimal contention and the maximum throughput. With the conservative¹ assumptions of these models, the maximum throughput of Noahnet is comparable to that of CSMA/CD network.

Section 2 of this paper gives an overview of the Noahnet flooding protocol for the sake of completeness. Section 3 gives the motivation for the performance modeling of Noahnet and derives a basic but important result for the throughput of Noahnet which is used in both models. Sections 4 and 5 each presents one performance model. Section 6 is the conclusion.

2 OVERVIEW OF NOAHNET FLOODING

Noahnet uses a graph-like topology with four to six interconnections per node. It uses flooding to route a data message from its source to its destination. The flooding of a message is achieved using three types of messages, namely, data, status, and command [1,3]. A data message is the one that

¹explained later in the paper why the assumptions of these models are conservative.



carries information. There are two types of status messages. The first type is used to indicate if the destination node has received a message with or without an error; the second type is used to indicate the flood status of a downstream node. A flood status can be one of "downstream", "blocked", or "got to the destination" (GTD). All the status messages are transmitted from a downstream node to its immediate upstream node. There is only one command message currently used which is "stop flooding." It is transmitted by an upstream node to its immediate downstream nodes indicating that the successful path has been established, and the downstream nodes should stop the flooding process immediately.

The routing of a message from a source to the destination node is a two step process - flooding-growth and flooding-contraction. During the growth of flooding, the message propagates to every node which is not occupied with a message and is reachable from the source node. During the contraction of flooding, the nodes that became occupied during the growth of flooding become unoccupied again. Nodes on unsuccessful paths become unoccupied in a relatively short time compared to the nodes on the successful path.

2.1 Flooding-Growth

In Noahnet flooding, whenever a node detects a start of a data message (SOM) from one of its adjacent nodes, it starts forwarding the message to all its unoccupied² adjacent nodes and starts sending occupied status to its occupied adjacent nodes. It should be noted that a node can be occupied with only one data message at any time.

The adjacent nodes again forward the message to their unoccupied adjacent nodes and start sending occupied status to their occupied adjacent nodes. This process continues until the message cannot be forwarded any more, that is, nodes do not find any unoccupied adjacent nodes to forward the message to when they detect start of this message.

The path of the message during the growth of flooding forms a rooted spanning tree of the network graph as shown in figure 1. The root of the tree represents the source node, and other nodes of the tree represent the nodes occupied by the same message. A message propagates from the root towards the leaf nodes of the tree.



2.2 Flooding-Contraction

During the spawning of a message tree (that is, the growth of flooding), the unoccupied nodes of the network become part of the tree. The tree consists of possibly one successful path and zero or more unsuccessful paths. The spawning of the tree is then followed by a contraction process which releases nodes on unsuccessful paths. The contraction process begins at leaf nodes and at the destination node, if a successful path exists as part of the tree. The leaf nodes start what is called the “blocking” process, whereas the destination node starts the “stop flooding” process. Figure 2 illustrates the “blocking” and “stop flooding” processes.

2.2.1 Blocking Process

A node is blocked if it is not the destination of the message it is occupied with, and either it does not have any downstream nodes, or all its downstream nodes also found themselves blocked. A blocked node sends a “blocked”

²Every node sends its status to all its adjacent nodes all the time except when it is transmitting a data message to the adjacent node.



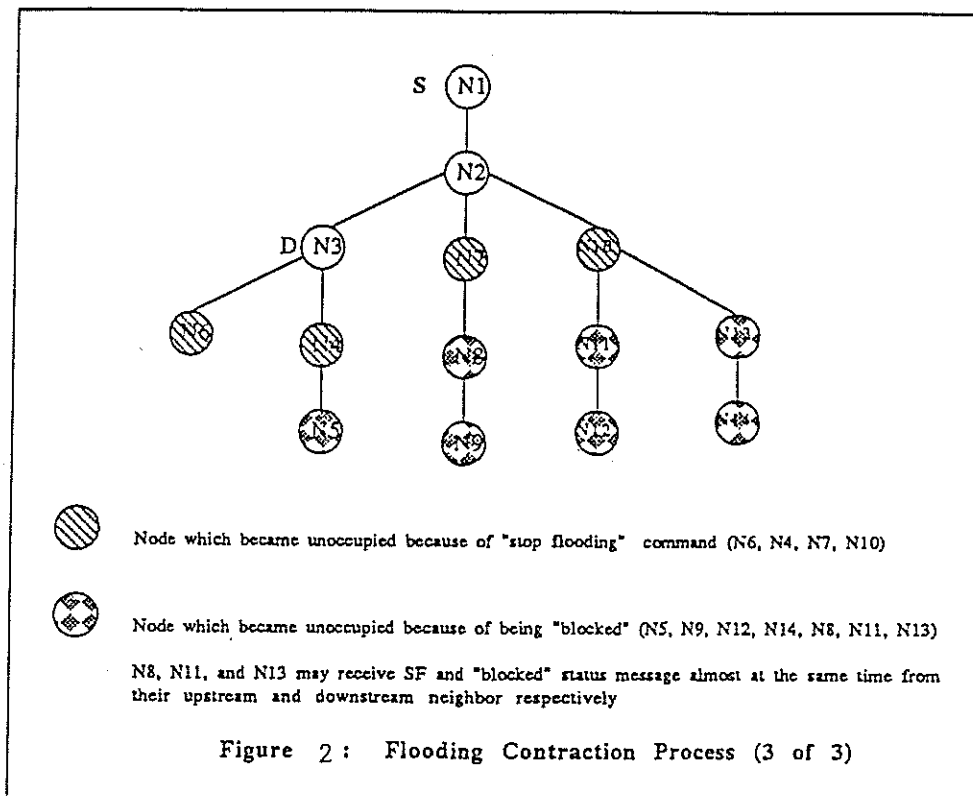


Figure 2: Contraction Process

status message to its upstream node, before getting ready to be unoccupied again.

Leaf nodes are the first to become blocked, and it is easy to see that the blocking process progresses from leaf nodes to upstream nodes, and it results in the freeing of nodes on unsuccessful paths.

2.2.2 “Stop Flooding” Process

The destination node is another place in the tree where the contraction process begins. As soon as the destination node detects that it is the destination of the message, it sends a premature End of Message (EOM) followed by a “stop flooding” (SF) command to all its downstream nodes, and GTD status message to its upstream node.

When a node receives an EOM followed by a SF command, it knows that a successful path has been already identified, and this node does not have to be occupied with this message any longer. v On the other side, when a node receives a GTD status message from one of its downstream nodes, it knows that it is on the successful path to the destination, and the successful path exists via the downstream node from which GTD is received. Thus, the paths via other downstream nodes constitute unsuccessful paths, and nodes on these unsuccessful paths can be released by sending an EOM followed by SF to the downstream nodes. Also, the node which received GTD status message sends GTD to its upstream node.

2.3 General Comments

It should be noted that as soon as a node becomes occupied with a message, it becomes a part of the tree being spawned by the message, and it can be either on the successful path or on an unsuccessful path. At any time, multiple messages can be active in the network. Each message spawns its own corresponding tree when it is looking for the destination. Thus, multiple simultaneous messages have the effect of dividing the network graph into partitions, where each partition is a tree spawned by the message being flooded in that partition.

If the destination node of a message is not in the partition of the message, the message is unsuccessful, and the source node finds itself blocked. In such a situation, the source node retransmits the message at a later time.

The time a node remains occupied as part of an unsuccessful path is called t_{us_path} , which may vary for each node on an unsuccessful path. The time a node remains occupied as a part of a successful path is called t_{s_path} . The ratio t_{us_path}/t_{s_path} is always less than one and is an important parameter which indicates how fast, compared to the transmission time of the message, nodes on unsuccessful paths become free. Reference [2] shows how to compute this ratio.

It should be noted that in a reasonable size Noahnet, a message can go to a node more than once, and thus, a message can loop in the network. Again, reference [2] looks at ways of avoiding this looping. The current approach is to keep the node occupied for an additional amount of time such that when a node becomes free, there is no node in the network looking for free adjacent nodes. This strategy is simple and effective but increases t_{us_path} .

3 MOTIVATION AND BACKGROUND

The Noahnet flooding protocol is fairly complex and Noahnet throughput and delay characteristics are far from obvious. The graph topology of Noahnet is inherently more reliable than bus and ring topologies of traditional LANs. It is also able to support multiple simultaneous dialogues. However, as the flooding creates unsuccessful paths, in addition to the possible successful path, it is not obvious how many simultaneous dialogues can be in progress in a network of N nodes. The purpose of these models is to get a better understanding of throughput characteristics of Noahnet. This may suggest modifications, trade-offs, or both in the existing Noahnet flooding protocol. These models also provide a basis for comparing Noahnet performance with that of other LANs such as Ethernet.

As with any other performance model, these models abstract the operation of Noahnet such that one does not have to worry about all the details, and can concentrate on the relevant issues of the protocol. The following assumptions attempt to abstract the important aspects of the protocol for the purpose of these models.

- Both models assume slotted Noahnet operation. This means that there are fixed time slots, and a node with a message to transmit attempts the transmission only at the beginning of the slot. The slotted as-

sumption considerably simplifies the analysis but introduces some approximations, because in reality, Noahnet does not operate in a slotted manner. However, simulation studies have shown that the approximations thus introduced are not significant, and the assumption does not change the pattern of results[3]. Also, the the experience with CSMA and CSMA/CD modeling shows that the models with slotted assumption can predict the performance quite accurately[4].

- The models assume very little about the network topology and the protocol³. It is assumed that if k messages are attempting transmission in a slot, the protocol and the topology are such that the network is divided into k partitions of arbitrary sizes; one partition for each message. A node can be in only one partition.

Each partition may have one successful path, if it exists, and zero or more unsuccessful paths for the message being flooded within that partition. Nodes on unsuccessful paths remain occupied for t_{us_path} , whereas nodes on successful path remain occupied for t_{s_path} . For the purpose of these models, it is assumed that $t_{s_path} = t_{us_path} = slot\ length$. In this sense, the model is pessimistic in predicting throughput, because in actual Noahnet, the nodes on unsuccessful paths will become unoccupied in a fraction of the message transmission time (that is, $t_{us_path} < t_{s_path}$) and can participate in the flooding or transmission of other messages in the remaining time. Thus, these these models predict some kind of a lower bound on the Noahnet throughput.

This assumption makes the model mostly independent of the protocol details. For example, the model does not depend on the value of the ratio of t_{us_path} and t_{s_path} or on how fast status and command messages propagate in the network. As long as the $t_{us_path} \leq t_{s_path}$, the model is valid.

- In addition, it is assumed that all the nodes are equiprobable destinations of the messages which originate from the other nodes.

For the purpose of these models, the throughput of the network is defined

³The advantage is that the models are quite independent of the minor details of the protocol and are useful for a family of protocols. The disadvantage however is that the models do not predict the effect of the various trade-offs in the protocol.

as the average number of messages successfully transmitted in a slot. Similarly, the load is defined as the average number of messages that attempted transmission in a slot.

In the rest of this section, an expression is derived for the throughput of the network which is used in the subsequent sections.

3.1 Average number of successful messages given k attempted transmissions

With the assumptions stated above, the following equation for the average number of successful messages in a slot given k attempted transmissions, γ_k , is derived.

$$\gamma_k = \frac{M - k}{N - 1} \quad \text{for } k \geq 1 \quad (1)$$

where

- M = Number of unoccupied nodes which participate in the flooding, $M \leq N$
- N = Total number of nodes in the network, including the nodes which do not participate in the flooding in the given slot.
- k = Number of messages that attempt transmission in this slot

Assuming that k nodes originate transmission of one message each, the unoccupied network of M nodes will be divided into k partitions; one partition for each message being flooded. A message would be successful, that is, the message would be successfully routed to its destination, if the destination belongs to the partition of the message.

It should be noted that

$$\gamma_k = 0 * P(0) + 1 * P(1) + \dots + k * P(k) \quad (2)$$

where $P(j)$ is the probability of j successful messages in the slot, $0 \leq j \leq k$.

Considering that all the nodes are equally probable as destinations for the messages originating at every node, the probability, p_i , that the i^{th} message is successful ($1 \leq i \leq k$) is the same as the probability that the destination belongs to the partition of the i^{th} message. Thus,

$$p_i = \frac{m_i - 1}{N - 1} \quad (3)$$

where m_i = the size of the partition of the i^{th} message. It should be noted that the partition also includes the source node which cannot be the destination, which is why 1 is subtracted from m_i . The sizes of the partitions satisfy the relationship $\sum_{i=1}^k m_i = M$.

Equation (1) is proved by using induction on k , the number of attempted messages in a given slot.

As the basis for the induction, consider $k = 1$ and $k = 2$. For $k = 1$, equation (2) reduces to

$$\begin{aligned} \gamma_1 &= 1 * P(1) \\ &= \frac{M - 1}{N - 1} \end{aligned}$$

This follows from the fact that $P(1) = p_1$, when there is only one partition, and $m_1 = M$.

For $k = 2$, equation (2) reduces to

$$\begin{aligned} \gamma_2 &= 1 * P(1) + 2 * P(2) \\ &= p_1 * (1 - p_2) + p_2 * (1 - p_1) + \\ &\quad 2 * p_1 * p_2 \end{aligned}$$

Substitution for p_i 's from equation (3) and simplification gives

$$\gamma_2 = \frac{M - 2}{N - 1}$$

Thus, the basis for the induction holds. As the induction hypothesis, assume that equation (1) holds for all i 's such that $i \leq k$. To complete the proof by induction, it is shown that equation (1) holds for $i = k + 1$ also.

Assume that there are $k + 1$ messages that attempt transmission in a slot. So the network gets divided into $k + 1$ partitions. Let the size of the partition for the i^{th} message be m_i .

Out of $k + 1$ partitions, consider only k partitions (say $m_1 \dots m_k$). Within these partitions, there are k attempted transmissions and $M - m_{k+1}$ nodes

that participated in the flooding. So the average number of successful messages out of these k messages, as given by the induction hypothesis, is equal to

$$\frac{M - m_{k+1} - k}{N - 1}.$$

Now, consider the m_{k+1} th partition. Again, using the induction hypothesis, the average number of successful messages in this partition is equal to

$$\frac{m_{k+1} - 1}{N - 1}.$$

So the total average number of successful messages given $k + 1$ attempted transmissions is given by :

$$\begin{aligned} \gamma_{(k+1)} &= \frac{M - m_{k+1} - k}{N - 1} + \frac{m_{k+1} - 1}{N - 1} \\ &= \frac{M - (k + 1)}{N - 1} \end{aligned} \quad (4)$$

Thus, equation (1) holds for $k + 1$ also, and therefore, for all $k \geq 1$.

3.2 Throughput

Let T_k denote the probability of k attempted transmissions, then the average number of successful messages, γ , is given by :

$$\gamma = \sum_{k=1}^M \gamma_k * T_k$$

Substituting for γ_k from equation (1) and simplifying gives

$$\gamma = \frac{M - \bar{k} - MT_0}{N - 1} \quad (5)$$

where \bar{k} is the average number of messages that attempt transmission in a slot.

This equation suggests that the average number of successful messages, which is the throughput, is a function of load (\bar{k}) and the probability of zero arrivals in a slot (T_0). This result is the basis for the derivation of load-throughput equation for the two models in sections 4 and 5.

Equations (1) and (5) are useful and fundamental results as they are derived with minimum assumptions regarding the protocol and topology, which makes them applicable to the entire family of Noahnet protocols.

In the next two sections, we present two performance models of Noahnet. Each model will have its assumptions, the throughput and load expressions, results in the form of plots, and the discussion of results.

4 MODEL 1 : with no retransmissions

4.1 Assumptions

This model makes the following assumptions regarding the Noahnet operation:

- Noahnet operates in a slotted manner and $t_{s_path} = t_{us_path} = \text{slot length}$ as assumed in the previous section.
- The input traffic also includes retransmissions. This means that if a message is not transmitted successfully, it disappears from the system and comes back as a fresh input message at a later time.
- The probability of a message arrival at each node in a slot is σ .
- All assumptions of the previous section are also in effect.

4.2 Load–Throughput

As all the assumptions of the previous section still hold, the average number of successful messages in a slot is given by equation (5). However, it should be noted that in every slot, all N nodes of the network participate in the flooding of messages, and therefore, we can substitute $M = N$ in the equation. Thus, the throughput, γ , is given by:

$$\gamma = \frac{N - \bar{k} - NT_0}{N - 1} \quad (6)$$

Since σ is the probability of a new arrival at a node, the average number of attempted transmissions per slot, \bar{k} , (the same as the load, λ) is equal to $N\sigma$. The probability of no arrival at any node in a slot, T_0 , is equal to $(1 - \sigma)^N$.

Substituting for \bar{k} and T_0 in the equation (6), gives

$$\gamma = \frac{N}{N - 1} \{1 - \sigma - (1 - \sigma)^N\} \quad (7)$$

The throughput in a slot, γ , can also be represented in terms of the load, ($\lambda = N\sigma$), as follows.

$$\gamma = \frac{N}{N - 1} \left\{1 - \frac{\lambda}{N} - \left(1 - \frac{\lambda}{N}\right)^N\right\} \quad (8)$$


4.3 Discussion

Equation (8) is plotted in figure 3 which shows how γ varies with λ for different values of N . Obviously, various values of λ are obtained by choosing values of σ between 0 and 1.

These results suggest that Noahnet is essentially a contention system where the throughput increases with load only up to a certain threshold value, beyond which the throughput starts decreasing with increasing load. The threshold value of load can be found by differentiating equation (8) with respect to λ , and equating it to zero, giving

$$\lambda_{max-\gamma} = N - N^{\frac{N-2}{N-1}}$$

These plots show that Noahnet is unstable in the sense that for a given value of throughput, the network can potentially operate at two values of the load. Also, there exists a large value of load which results in zero throughput, because at that load, the contention is so much that no message is transmitted successfully. For example, consider $\lambda = N$, that is, every node in the network is trying to send its own message. As every node has a message to send, no node can receive a message, and thus, the number of successful messages is zero (or throughput equal to zero) which is what model correctly predicts.



3

Figure 4: Model 1 : Load-Throughput Behavior

The maximum throughput, that is the maximum average number of messages which can be transmitted in a slot is less than or equal to 1. Figure 3 shows that the threshold value of load and the maximum throughput increase with number of nodes in the network. In fact, taking limit as $N \rightarrow \infty$ in equation (8) gives

$$\gamma = (1 - e^{-\lambda})$$

which is also shown in figure 3. This figure shows that for large number of nodes in the network, the throughput does not drop with increased load.

The most important conclusion of this model is that the graph topology by itself does not guarantee that the network can have multiple successful messages in a slot to give higher throughput or to better support multiple burst dialogues. This means that the protocol and the topology has to be designed carefully to achieve higher throughput from a network with the graph topology.

5 MODEL 2 : with retransmissions

5.1 Assumptions

This model is based on the same assumptions as the first model except the retransmission strategy. This is modeled as follows:

- The input traffic does not include retransmissions. If a message is not transmitted successfully, it becomes a back-logged message at its source node, and is retransmitted in the next slot with a certain probability α by its source node. This may be contrasted with model 1 where the unsuccessful message disappears from the system and comes back as a new message at a later time, so that, the input traffic also includes the retransmissions and is equal to the load. In model 2, the input traffic equals the throughput if the network is stable.
- A node with a backlogged message transmits the message in the subsequent slot with probability α .
- The probability of a node being backlogged is β .

- The probability of a message arrival at a node is σ . Only nodes with no back logged messages can receive a new arrival. This means that a node can have at most one back-logged message.

5.2 Load–Throughput Expressions

We know that the average number of successful messages in a slot, γ , can be obtained from equation (5) with $M = N$. Thus, γ is given by:

$$\gamma = \frac{N - \bar{k} - NT_0}{N - 1} \quad (9)$$

where

- T_0 is the probability of no transmission in a given slot. No transmission in a slot for this model means that no new arrivals take place and none of the back-logged nodes attempt retransmission in the slot. Thus, T_0 is given by:

$$T_0 = \{1 - \sigma - \beta(\alpha - \sigma)\}^N \quad (10)$$

- \bar{k} is the average number of attempted transmissions in a slot, which in this model includes new arrivals and retransmissions, and is given by:

$$\bar{k} = N\{(1 - \beta)\sigma + \beta\alpha\} \quad (11)$$

Comparing equations (10) and (11), we get $T_0 = \left(\frac{N - \bar{k}}{N}\right)^N$. Substituting this value of T_0 in equation (9) gives the following result.

$$\gamma = \frac{N - \bar{k} - N\left(\frac{N - \bar{k}}{N}\right)^N}{N - 1} \quad (12)$$

where \bar{k} is the load, that is, the average number of attempted transmissions in a slot.

This is an interesting result because the throughput depends only on the load and does not directly depend on other parameters, such as α, β , or σ .

Lastly, we also know that if the network is stable, the throughput is equal to the average input rate. Thus, γ is also given by

$$\gamma = (N - N\beta)\sigma \quad (13)$$

5.3 Discussion

It should be noted that equation (12) of model two is the same as equation (8) of model one with $\bar{k} = \lambda$. These equations suggest that for the same number of attempted transmissions, the average number of successful messages is the same, and it does not matter if the load consists of all new arrivals or a mix of new arrivals and retransmissions. Thus, as far as the load-throughput behavior is concerned, the comments from model one carry over to this model.

However, the purpose of this model is to understand the effect of the retransmission probability, α , on the throughput of the network. This is achieved by choosing various values of σ and α in equations (12)⁴ and (13) and solving for β to get the throughput from equation (12) or (13)⁵.

Figure 4 shows how the throughput varies with α for different values of σ and N . The following observations can be made from these figures.

- For small values of σ ($\sigma \leq 0.06$ for $N = 16$), that is, for a small number of new arrivals in a slot, the throughput is almost independent of α , the retransmission probability. The small number of new arrivals means that there is little contention in the network, and therefore, the nodes with backlogged messages can afford to retransmit their messages often.
- For moderate values of σ ($0.06 < \sigma < 0.3$), the throughput remains high only up to a certain value of α . Beyond this threshold value, the retransmission attempts increase the contention so much that the throughput starts decreasing rapidly. And obviously, the network should not operate in this range of α and σ .
- For large values of σ ($\sigma > 0.3$), that is, for a large number of new arrivals in a slot, the throughput rapidly decreases with α beyond even a small value of α .
- For $\alpha = 1$, the steady state throughput is 0 for all σ 's. This can be explained as follows.

The value of α equal to one implies that a node with a back-logged messages will attempt retransmission of its messages in every subsequent

⁴Note that \bar{k} is given by equation (11) in terms of α, β , and σ .

⁵It should be obvious that the probability of a node being blocked, β , cannot be selected independently as all three variables σ , α , and β cannot be independent.

Throughput : Successful M's in one slot




Figure 4: Model 2 : Load-Throughput with alpha (1 of 2)

Throughput : Successful M's in one slot




Figure 4: Model 2 : Load-Throughput with alpha (2 of 2)

slot. And there is a finite probability that the number of back-logged nodes is equal to N . So when all nodes become back-logged, all N nodes would try to transmit a message in every slot with probability of success of a message equal to zero. Therefore, once all nodes are back-logged, the system stays in this state for ever and no message ever gets transmitted successfully. Thus, α must not be equal to 1 to achieve nonzero steady state throughput.

In short, the values of σ and α should be chosen such that the contention in the system is optimal - not too much to give less throughput and not too little to underutilize the resources and give less throughput. The plots suggest the possible range of values which give satisfactory results.

6 CONCLUSION

We have presented two preliminary performance models for an experimental flood local area network, Noahnet. The purpose of these models was to give us a better understanding of load-throughput characteristics of Noahnet. However, these models assume very little about the protocol and the topology. Therefore, these models are fairly general purpose, and can be used for other similar systems not necessarily using the same flooding protocol.

Another important point to note about these models is that they give a lower bound on Noahnet throughput because they assume that nodes on the unsuccessful paths remain occupied for the duration of the message transmission time. In actual reality though, the nodes on unsuccessful paths remain occupied only for a fraction of that time, and in the rest of the time, they participate in the flooding of other messages to give higher throughput.

We can summarize the conclusions as follows.

- The maximum throughput of the network is always less than one message per slot (or per one transmission time of a message). This means that the graph topology by itself is not adequate to give multiple successful transmissions in a slot. The protocol and the topology have to be designed carefully to achieve higher throughput.
- The Noahnet is unstable in the sense that the throughput increases with the load only up to a threshold value of load, beyond which it starts

decreasing with the load. As a result, for a given value of throughput, the network can operate at two possible loads.

- Noahnet is a contention system, and a message is not necessarily transmitted successfully in one attempt. Thus a node sometimes has to retransmit a message more than once. The second model suggests that for a given value of σ (measure of new arrivals), value of α (measure of retransmissions) should be chosen such that the contention is optimal. In other words, the contention should not be so little as to underutilize the resources and should not be so high that no successful transmissions take place.
- The throughput of Noahnet under the pessimistic assumptions, as predicted by these models, is comparable to CSMA/CD networks. The reason being that the maximum throughput of CSMA/CD networks is also always less than one message in a slot, and at very large loads, the throughput tends to decrease with the load because of the increased contention.

It should be noted however that Noahnet does require more number of communication links, and therefore, Noahnet is more expensive than CSMA/CD networks.

Work is in progress on two more performance models of Noahnet which will supplement the models presented in this paper. One of these new models attempts to provide the load-delay characteristics of Noahnet under the similar assumptions. The other model attempts to take into account more details of the protocol. More specifically, this model takes the time the nodes on unsuccessful paths remain occupied and the connectivity of the network as parameters of the model.

Acknowledgement We will like to thank Chuck Cranor for the time he spent proof reading early drafts of this paper.

References

- [1] Farber, D.J., Parulkar, G.M., *Noahnet: An Experimental Flood Local Area Network*, Proceedings of Eighth International Conference on Com-

puter Communication (ICCC), pp: 265-269, Munich, West Germany, Sept 1986.

- [2] Farber, D.J., Parulkar, G.M., *A Closer Look at Noahnet*, Proceedings of ACM SIGCOMM'86 Symposium on Communications Architectures and Protocols, pp: 205-213, Stowe, Vermont, Aug 1986.
- [3] Parulkar, G.M., *Noahnet: An Experimental Flood Local Area Network*, Ph.D. Thesis, Department of Computer Science, University of Delaware, 1987.
- [4] Hammond, J.L., O'Reily, P., *Performance Analysis of Local Computer Networks*, Addison-Wesley Publishing Company, 1986.