An Analysis of TCP over Random Access Satellite Links

Chunmei Liu and Eytan Modiano
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: mayliu, modiano@mit.edu

Abstract— **This paper analyzes the performance of TCP over satellite links with random access. The system considered consists of large number of identical source-destination pairs, each employing TCP at its transport layer and a random access scheme at its MAC layer. Simple formulas that well capture the system performance are given, and some important properties of the system performance are presented as well. Specifically, in order to analyze the system, we propose a simplified version of the system. We then develop formulas for obtaining the throughput of this simplified system as a function of various system and protocol parameters. Based on these formulas, it is shown that the maximum possible system throughput is $1/e$, which can be achieved only when the system parameters satisfy a given condition. The optimal MAC layer transmission probability at which the throughput is maximized is derived as well. Furthermore, the impact of varying system and protocol parameters on the system performance is analyzed. The results show that for systems with very small propagation delay or very large number of source-destination pairs, a throughput of $1/e$ can be achieved by setting the MAC layer transmission probability to its optimal value. However, when the number of users is (relatively) small or the propagation delay is (relatively) large, the maximum achievable throughput can be substantially smaller than $1/e$. Although the analysis is based on the simplified system, simulations on the original system show that the formulas and the above results can be used to describe the performance of the original system as well.**

I. Introduction

In this paper we consider the performance of TCP over a shared satellite link with random access (i.e., Aloha). Satellite links are inherently different from terrestrial links, and the performance of TCP over satellite has received a great deal of attention over the past decade. Most previous work has focused on the impact of the large propagation delays and high bit error rates associated with satellite links, [1], [4], [5], [6]. Overviews are given in [3], [8].

This paper addresses the multi-access nature of satellite networks, and analyzes the TCP performance over satellite links when a pure random access scheme is employed at the MAC layer. In particular, the system considered consists of large number of identical persistent source-destination pairs, each pair employing TCP at its transport layer and each source employing a random access scheme at its MAC layer. In order to analyze the system, we propose an simplified system that differs from the above system in that the transport layers employ a simplified version of TCP window flow control. For this simplified system, an analytical model is developed, and equations are derived for solving for the system throughput, idle probability and collision probability. Based on the equations, we show that the maximum achievable throughput is $1/e$ and give a necessary and sufficient condition for achieving it. The impact of system and protocol parameters on the system performance is obtained both analytically and numerically. Finally, the performance of the original system is examined by simulations and the results are compared with those numerical results obtained from the formulas for the simplified system. The results show that although the analysis and formulas are based on the simplified system, they are a good approximation of the original system and all the results obtained are valid for the original system as well.

The paper is organized as follows: the next section gives a detail description on the systems and protocols examined. In Section III we develop an analytical model for the simplified system. Using this analytical model, Section IV gives formulas for the system's throughput, the conditions for achieving the maximum throughput, as well as the optimal MAC layer transmission probability. The impact of various system parameters on performance is also examined in Section IV, both analytically and numerically, and simulation results for the original system are presented and compared with the numerical results as well. Section V discusses the implications of the various assumptions and concludes the paper.

II. System Description

The system we consider consists of N identical source-destination pairs (SD pairs), where N is large, as shown in Figure 1. Each of the N sources has unlimited number of packets to be sent to its corresponding destination. All N sources share a common channel to the satellite. Each SD pair employs TCP at its transport layer for congestion control purpose, and all sources employ a simplified version of ALOHA multi-access scheme at their MAC layer for multi-access purpose. For simplicity, we ignore the timeout update of TCP based on its round trip time measurements, and assume that the timeout value is a random variable with mean TO slots. The ALOHA schemes will be described in detail later.

For all SD pairs, all packets have the same length and each packet requires one time unit (called a slot) for transmission. The two-way propagation delay, defined to be the duration between the time when a packet is successfully transmitted by the MAC layer of the source and the time when its corresponding acknowledgement is received by the source, is a random variable with mean D slots. Here the randomness represents the queuing delays and other random delays experienced by the packets. To focus on the impact of random access, assume that there are no other packet losses in the network except losses due to MAC layer collisions.

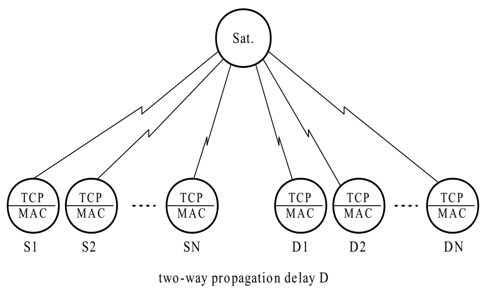

Fig. 1. System with TCP over MAC Random Access

The simplified version of ALOHA multi-access scheme works as follows:

- When one or more packets are available in the MAC layer buffer, the MAC layer transmits the first packet with probability p.
- If two or more packets are transmitted in a given slot, a collision occurs and all packets are assumed lost. If only one packet is transmitted in a given slot, the packet is successfully received.

The difference between this scheme and ordinary ALOHA lies in that here we do not allow retransmissions of packets involved in collisions at the MAC layer. This assumption is reasonable over a satellite link as retransmissions at the MAC layer will almost certainly result in a TCP time-out event, which can be easily verified. As with ALOHA, each MAC layer of the system can be viewed as a queuing system, with all arrivals coming from its corresponding transport layer and geometrically distributed service time with mean $1/p$.

In order to obtain an analytical model, we consider a simplified system that differs from the above system only in that it has a different window flow control (WFC) scheme. Its WFC scheme is similar to that in the congestion avoidance phase of TCP but without the fast retransmit/recovery option. In particular, it works as follows:

- Initially the window size $W = 1$.
- Upon an acknowledgement, the window size is updated to be $W \leftarrow W + \frac{1}{\lceil W \rceil}$.
- Upon a timeout, the window size is reset to be $W \leftarrow 1$.

We also assume that the transport layer employs a selective repeat retransmission protocol and retransmits only those packets that were lost. Furthermore, we assume that the timeout value is large enough that the probability of timeout due to queuing delay at the MAC layer can be ignored. Recall that there are no losses other than collision losses. Hence, timeouts can only be triggered by MAC layer collisions.

TABLE I
WINDOW EVOLUTION AND ROUNDS

ACK	window size	packet released	round	first packet in round	number of packets in round
	1	1	1	1	1
1	2	2,3	2	2	2
2	2	4	3	4	3
3	3	5,6			
4	3	7	4	7	4
5	3	8			
6	4	9,10			
7	4	11	5	11	5
8	4	12			

The analysis henceforth will be based on the simplified system with these assumptions, and we therefore call the simplified system the system when there is no ambiguity. Nevertheless, the analytical results match well with the simulations for the original system with TCP without timeout update, as will be presented in Section IV. Including timeout update into analysis is part of our future work.

For convenience, we borrow the concept of round from [7] at the transport layer, with extension that includes the timeout signal: a round starts with transmission of W packets, where W is the current window size of the congestion window. Once all packets falling within the congestion window have been sent, no other packets are sent until the first ACK for one of these W packets or a timeout signal is received . This ACK reception or timeout marks the end of the current round and the beginning of the next round.

We further index the packets sent and the rounds between two successive timeouts in order, i.e., packet k is the kth packet and round k is the kth round. The ACK for packet k is called ACK k. Table I illustrates the window evolution and rounds between two successive timeouts for the simplified system. For example in Table I, each time when the window size is increased by 1, there are two packets released. Otherwise, there is one packet released. After the window size reaches 3, there are three ACKs received, namely ACK 3,4 and 5, before the window size reaches 4. Correspondingly, four packets are released, namely packet 5, 6, 7 and 8. These are shown in column 1 to column 3. Round 4 begins with ACK 4, which is the first ACK of packets in round 3, and the first packet in round 4 is packet 7. ACK 7 thus becomes the first ACK of packets in round 4 and marks the end of round 4 and beginning of round 5. This is shown in the 4th and 5th columns. By counting the number of packets in round 4, we obtain 4, which is shown in the last column.

III. PERFORMANCE ANALYSIS

In the following we first consider the entire system including all SD pairs, and then examine one SD pair in isolation. The combined analysis yields formulas that together can be used to solve for the system performance.

A. Analysis at the System Level

For every sender/receiver pair, define the throughput as the number of packets correctly received by the receiver per unit

0-7803-8344-3/04/$20.00 (C) 2004 IEEE

time, and the send rate as the number of packets sent by the sender per unit time. That is, the send rate includes both packets that are included in throughput and the packets that are lost due to collisions. By assuming equivalence between time average and ensemble average, the throughput is also the probability that in a slot, a packet is correctly received by the receiver. Similarly, the send rate is also the probability that in a slot, a packet is sent by the sender.

In our system, on one hand, we have two layers, the transport layer and the MAC layer. We therefore have definitions of throughput and send rate at both layers. On the other hand, the MAC layers receive packets only from their corresponding transport layers. Therefore, the send rate at both layers are the same. In addition, collision losses at the MAC layers are the only losses in the system. So the throughput at both layers are the same as well. Henceforth, we use the name throughput and send rate to refer to the throughput and send rate at both layers. Besides, the throughput and send rate can refer to those of a single SD pairs and those of all N SD pairs. We call those for a single SD pair individual throughput and send rate, and denote them by λ_d and B_d, respectively. We call those for all N SD pairs system throughput and send rate, and denote them by λ_s and B_s, respectively.

From the above definitions and discussions of throughput and send rate, we conclude that in our system, the system throughput λ_s is the probability that there is exactly one MAC layer sending a packet in a slot. For one particular SD pair, the individual send rate B_d is the probability that the MAC layer of this SD pair sends a packet in a slot. Since we have a large number of SD pairs (N is large), we further assume that for a particular SD pair, its state is independent of the state of other SD pairs. We hence have the following relationship between the system throughput, idle probability, collision probability and the individual send rate:

$$\lambda_s = NB_d(1 - B_d)^{N-1} \approx NB_d e^{-NB_d}, \tag{1a}$$
$$P_I = (1 - B_d)^N \approx e^{-NB_d}, \tag{1b}$$
$$P_C = 1 - \lambda_s - P_I. \tag{1c}$$

where P_I and P_C are the idle probability and collision probability, respectively, and the approximations hold when B_d is small and N is large.

On one hand, due to the independent assumption, the above analysis is similar to that of a standard ALOHA system. The equations in (1) show that, as in an ALOHA system, the number of packets correctly received in a slot can be approximated by a Poison random variable, with the attempt rate NB_d as its mean. The maximum possible throughput can be achieved at $NB_d = 1$, with the corresponding throughput, idle probability and collision probability being $1/e$, $1/e$ and $1 - 2/e$, respectively. On the other hand, different from an ALOHA system, $NB_d = 1$ is not always achievable in our system due to the transport layer window limitation, which will be analyzed in detail later, while in an ALOHA system, attempt rate 1 can always be achieved with proper parameter settings. We henceforth call the system performance at $NB_d = 1$ ALOHA performance.

Furthermore, under the independent assumption, the effects of other SD pairs on one particular SD pair are aggregated into one parameter Q, defined to be the probability that all other $N - 1$ MAC layers transmit no packets in a slot. Since all SD pairs are identical, we have:

$$Q = (1 - B_d)^{N-1} \approx e^{-(N-1)B_d} \approx e^{-NB_d}. \tag{2}$$

Note that although the values of P_I and Q are approximately the same (Equation (1b) and (2)), they have different physical meaning. P_I is for all N SD pairs, while Q excludes the SD pair considered and is for the other $N - 1$ SD pairs only.

Moreover, the independent assumption also gives $\lambda_s = N\lambda_d$ and $B_s = NB_d$.

B. Analysis at the session level

As mentioned before, the effects of the other $N-1$ SD pairs on one particular SD pair are aggregated into one parameter Q, the idle probability of all other $N-1$ SD pairs. This particular SD pair can thus be modelled as a normal transport layer session with collision probability of each packet being $1 - Q$. This section models one particular SD pair as a renewal process and obtain an upper bound and a lower bound for its send rate, B_d, in terms of Q.

For one particular SD pair, let's consider after a timeout, what happens before the first collision occurs. All packets sent before the collided packet won't encounter a timeout, since they encounter no losses and the large timeout value assumption ensures that their ACKs will be received before the retransmission timer expires. Whereas the collided packet will eventually encounter a timeout, since there is no mechanism other than the transport layer retransmission to recover this collision loss. Therefore, the first collision after a timeout causes a successive timeout. We introduce the concept of cycle to denote the interval between two successive timeouts.

Moreover, since the transmissions of the packets between the first timeout and the collided packet were successful, the transport layer sender will receive some ACKs after the collision and release some new packets. The large timeout value assumption ensures that the MAC layer will finish the transmissions of these later released packets before the second timeout signal. Thus upon each timeout signal, there is no packet in the MAC buffer. In addition, according to the WFC scheme, the window evolves exactly the same after each timeout signal. The timeout signal sequence therefore forms a renewal process, and a cycle between two successive timeouts is an inter-arrival period of the renewal process. Let M be the number of packets sent during a cycle and T be cycle length. Then by the renewal theory,

$$B_d = \frac{E[M]}{E[T]}. \tag{3}$$

To solve for B_d, let's first consider $E[T]$. Instead of deriving its exact expression, which requires complicate queuing analysis, we derive an upper bound and an lower bound. Let R be the number of successful round during a cycle and RTO be the timeout value that ends the cycle. Then T is the sum of the duration of these R rounds plus RTO. Each round takes at least the service time of its first packet, denoted by F_k for round k, plus the two way propagation delay this packet experiences, denoted by D_k, for the sender to receive its ACK. We therefore have a lower bound for T, denoted by T^L, to be $T \geq T^L = \sum_{k=1}^{R}(F_k + D_k) + RTO$.

Similarly, an upper bound for T, denoted by T^U, can be obtained when ignoring the overlapping between the service time of each packet and D_k of each round k. Mathematically, define Z to be the index of the packet that incurs the first collision and X_k to be the service time of packet k. Then $T \leq T^U = \sum_{k=1}^{Z} X_k + \sum_{k=1}^{R} D_k + RTO$. Along with the lower bound T_L, we have

$$\sum_{k=1}^{R}(F_k + D_k) + RTO = T^L \leq T \leq T^U = \sum_{k=1}^{Z} X_k + \sum_{k=1}^{R} D_k + RTO.$$

By our multi-access scheme, the service time for each packet is geometrically distributed with mean $1/p$ and independent of each other and R and Z as well. Thus by taking expectations of the above inequality,

$$(\frac{1}{p}+D)E[R]+TO = E[T^L] \leq E[T] \leq E[T^U] = \frac{E[Z]}{p}+DE[R]+TO. \quad (4)$$

Now let's consider $E[M]$. Recall that packet Z is the packet that incurs the first collision. This means that the previous $Z-1$ packets were successfully transmitted and have been or will be ACKed, while packet Z and thereafter won't be ACKed. According to the WFC scheme, new packets can be released only upon reception of ACKs. Therefore, M equals to 1, counting for the first packet, plus the number of packets triggered by the $Z-1$ ACKs. In addition in the WFC scheme, each ACK triggers the release of one packet, except those ACKs that increase the window size by 1 (i.e., the last ACK in a round) where two packets are released (See Table I for illustration). Let I be the number of such ACKs in the $Z-1$ ACKs, then $M = 1 + (Z-1) + I = Z + I$. Moreover, the window size is increased by 1 per round. Therefore $I = R$, and

$$E[M] = E[Z] + E[R]. \quad (5)$$

We now have bounds for $E[T]$ as in Inequality (4) and $E[M]$ as in Equation (5) in terms of $E[Z]$ and $E[R]$. Recall that Q is the probability that no other senders send their packets in one slot and this event is independent of the state of the particular SD pair. Therefore, each packet of the particular SD pair incurs a collision with probability of $1-Q$ and independent of each other. Z is thus geometrically distributed with $Pr(Z=z) = Q^{z-1}(1-Q)$ and $E[Z] = 1/(1-Q)$. By exploring the relationship between R and Z, it can be shown that $E[R] = \sum_{k=1}^{\infty} Q^{k(k+1)/2}$. For brevity, we omit the details. By combining them with Equation (3), Inequality (4) and Equation (5), we obtain the following bounds for the send rate B_d:

$$\frac{\frac{1}{1-Q} + \sum_{k=1}^{\infty} Q^{k(k+1)/2}}{\frac{1}{p(1-Q)} + D\sum_{k=1}^{\infty} Q^{k(k+1)/2} + TO} = f^L \leq B_d$$
$$\leq f^U = \frac{\frac{1}{1-Q} + \sum_{k=1}^{\infty} Q^{k(k+1)/2}}{(\frac{1}{p}+D)\sum_{k=1}^{\infty} Q^{k(k+1)/2} + TO}. \quad (6)$$

Notice that for large enough timeout value TO, both bounds are increasing functions of Q. Mathematically, this can be shown by taking the derivative of the bounds with respect to Q. Physically, from our derivation, the lower bound corresponds to the send rate of the following system: after sending the first packet of each round, the MAC layer holds the transmission of other packets until it receives the ACK of the first packet. This is how we obtain the upper bound for the duration T between two successive timeouts. Clearly, the send rate of this system increases with the idle probability of the other $N-1$ SD pairs Q. Similarly, the upper bound corresponds to the send rate of the following system: after sending the first packet of each round, the MAC layer finishes the transmission of all other outstanding packets before it receives the ACK of the first packet, that is, within time D_k for round k. This is how we obtain the lower bound for the duration T between two successive timeouts. Again, the send rate of this system increases with the idle probability of the other $N-1$ SD pairs Q.

To fully under the system behavior, the expectation of the collision window W_z, defined to the window size when the collided packet Z is released (i.e., when the collision occurs), can also be shown to be:

$$E[W_Z] = 1 + \sum_{k=2}^{\infty} Q^{\frac{k(k+1)}{2}-2}. \quad (7)$$

Again for brevity, we omit the details.

IV. System Performance

Section III-A analyzes N SD pairs together and gives one relationship between the individual send rate B_d and the idle probability of $N-1$ SD pairs Q, in Equation (2). Section III-B gives an upper bound and a lower bound of B_d in terms of Q in Inequality (6). Based on these results, this section first gives bounds for B_d and Q in terms of system and protocol parameters, and then discusses the system performance. A sufficient and necessary condition for the system to achieve the ALOHA performance is also given, as well as the optimal MAC layer transmission probability at which the throughput is maximized. The analysis is confirmed by the numerical results under different system and protocol parameters. This section also gives the simulation results on the original system and compares them with the numerical results for the simplified system.

Before proceeding, let's first introduce how we set the system and protocol parameters in the numerical computations and simulations. For the systems, we consider the following data that is typical for a satellite link: the two-way propagation delay of one transmission is around 1sec; packet size L is about 10000 bits; and transmission rate is around 100k - 1M bps. Converting these into the parameters in our system, we obtain that the two-way propagation delay is about 10 - 100 time slots. We therefore set the two-way propagation delay of packets to be a uniformly distributed random variable with mean D being 10-100 slots and range of 4 slots. As mentioned before, the randomness represents the queuing delay and other random delayed experienced by packets.

The transport layer timeout value is also set to be a uniformly distributed random variable with mean TO. Similar to the timeout update in TCP, TO is set to be the average round trip time seen by the transport layer packets, denoted by RTT, plus four times its standard deviation. Because of the large propagation delay of satellite links, when p is not close to zero, we can ignore the queuing delays the packets experiences at the MAC layer and approximate RTT by the sum of packet service time at the MAC layer and the two-way propagation delay. Recall that the packet service time at the MAC layer is geometrically distributed with mean $1/p$ and variance $(1-p)/p^2$. Therefore, we obtain $TO \approx \frac{1}{p} + D + 4\frac{\sqrt{1-p}}{p}$. Notice that this setting of TO is used only in the numerical computations and simulations. The following analysis of impact of parameters on the system performance does not depend on this particular timeout value setting. It only requires that TO decreases with p and increases with D, which should be true for other TO settings as well.

Simulations are performed on ns simulator for the original system. The timeout value is set in the same way described above.

A. System Performance and Condition for ALOHA Performance

Based on the Equation (2) and Inequality (6), this subsection shows how to obtain the system performance given the system and protocol parameters.

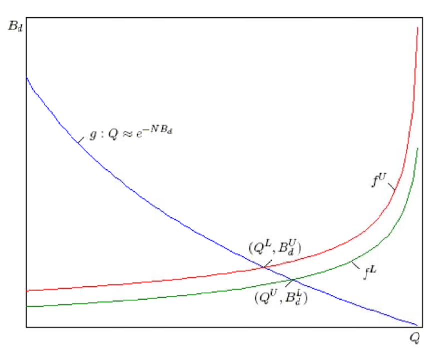

Fig. 2. Bounds for B_d and Q

Figure 2 plots the relationship between B_d and Q as given in Equation (2), called curve g. Figure 2 also plots the bounds of B_d, f^L and f^U, as given in Inequality (6). For clarity, the distances between the curves are exaggerated. The actual B_d and Q should be on curve g, as well as be inside the area between curve f^L and f^U. Therefore, B_d and Q are on the section of curve g between curve f^L and f^U. The intersections of curve g with curve curve f^L and f^U thus give a lower bound and an upper bound for B_d as well as a lower bound and an upper bound for Q. Denote them by B_d^L and B_d^U, Q^L and Q^U, respectively. Clearly, B_d^L and B_d^U are solutions of the following two equations, respectively:

$$B_d = \frac{\frac{1}{1-e^{-NB_d}} + \sum_{k=1}^{\infty} e^{-NB_d k(k+1)/2}}{\frac{1}{p(1-e^{-NB_d})} + D\sum_{k=1}^{\infty} e^{-NB_d k(k+1)/2} + TO}, \quad (8a)$$

$$B_d = \frac{\frac{1}{1-e^{-NB_d}} + \sum_{k=1}^{\infty} e^{-NB_d k(k+1)/2}}{(\frac{1}{p} + D)\sum_{k=1}^{\infty} e^{-NB_d k(k+1)/2} + TO}. \quad (8b)$$

Figure 3 plots the system performance, the system throughput, idle probability and collision probability, as a function of B_d, given in Equation (1). From the figure we can see that the bounds for B_d actually give us the range of the system performance. Specifically, if both bounds for B_d, B_d^L and B_d^U, lie within the same monotonic region of the throughput, i.e., within $[0, 1/N]$ or $[1/N, \infty]$ as in Figure 3, then their corresponding throughput, λ_s^L and λ_s^U in Figure 3, are also lower and upper bounds for the actual throughput λ_s. Otherwise, $B_d^L \in [0, 1/N]$ and $B_d^U \in [1/N, \infty]$, the actual throughput is close to the maximum possible throughput $1/e$. Furthermore, since P_C and P_I are monotonic with B_d, the collision probability and idle probability corresponding to B_d^L and B_d^U are also bounds for P_C and P_I. As in Figure 3, $P_C^L \leq P_C \leq P_C^U$ and $P_I^L \leq P_I \leq P_I^U$. We thus conclude that the two bounds given in Equations (8), together with Equation (1), fully characterize the system performance.

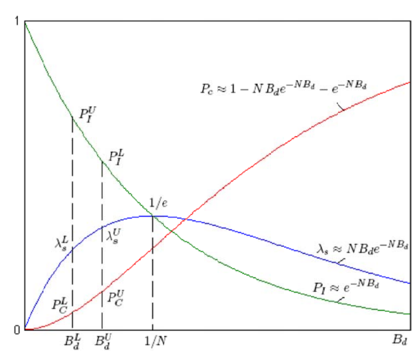

Fig. 3. System Performance with B_d

Our numerical results further show that in almost all cases, these two bounds are very close to each other. Table II gives some numerical results for different system and protocol parameters N, D and p. The difference between B_d^L and B_d^U is shown in the last column. Henceforth we approximate the actual B_d by B_d^U and use B_d and B_d^U exchangeably. In conclusion, the system performance is well approximated by Equations (8b) and (1).

Now let us find a sufficient and necessary condition on which the ALOHA performance, with the system throughput

TABLE II
DIFFERENCES BETWEEN B_d^U AND B_d^L

D	N	p	B_d^U	B_d^L	$B_d^U - B_d^L$
20	20	0.8	0.0528	0.0509	0.0019
30	10	0.9	0.0539	0.0524	0.0015
50	40	0.7	0.0254	0.0249	0.0005

$1/e$, can be achieved. Recall that this performance can be achieved if and only if $NB_d = 1$. By plugging $NB_d = 1$ into Equation (8b) and noting that its solution for B_d is unique (see Figure 2), we obtain the following sufficient and necessary condition for the ALOHA performance:

$$2.0N - 0.42D - TO - \frac{0.42}{p} = 0. \quad (9)$$

Notice that TO is a function of D and p. Also notice that due to the limit range of the actual parameters, the above condition cannot be always satisfied by adjusting one parameter with others fixed. For example for very large D, the solution of the above condition for p can be negative, while the actual transmission probability p has to be nonnegative. In this case, the ALOHA performance cannot be achieved by adjusting p only.

Overall, given the system and protocol parameters, the system performance can be solved from Equation (1) and (8b). Condition (9) is a sufficient and necessary condition for the system performance to achieve the ALOHA performance. Due to the limited range of system and protocol parameters, this condition cannot be always satisfied by adjusting one parameter with others fixed. That is, the ALOHA performance is not always achievable.

The following subsections analyze how the system performance changes with different system and protocol parameters.

B. Impact of Transmission Probability on System Performance

First consider the impact of transmission probability p on the system performance. By the definition of f^U (Inequality (6)) and noting that TO is a decreasing function of p, f^U increases with increasing p. Thus as p increases, curve f^U in Figure 2 moves up. On the other hand, curve g is not a function of p and remains the same. Therefore, the intersection of curve g and f^U moves leftwards. Consequently, B_d^U increases monotonically with p.

Recall that condition (9) is a sufficient and necessary condition for B_d^U to achieve $1/N$. Also note that $p \in [0, 1]$ and when $p = 0$, $B_d^U = 0$. Therefore, when p increases from 0 to 1, if the solution of condition (9) for p, denoted by p_{max}, lies within $[0, 1]$, then B_d^U increases from 0 to $1/N$ then to some number. Consequently, the throughput first increases, then decreases, with maximum $1/e$ (see Figure 3). Otherwise, B_d^U increases from 0 to some number below $1/N$. Consequently, the throughput increases monotonically with p, and the maximum throughput is achieved at $p = 1$.

Similarly, since P_I and P_C is monotonic with B_d, P_I decreases monotonically with p, and P_C increases monotonically with p.

The above discussion actually gives us the transmission probability at which the throughput achieves its highest value, denoted by p_{opt}, as follows:

$$p_{opt} = \begin{cases} p_{max} & \text{when } p_{max} \in [0, 1] \\ 1 & \text{otherwise} \end{cases} \quad (10)$$

Moreover, if $p_{opt} = p_{max} \in [0, 1]$, then the system achieves ALOHA performance. Otherwise, $p_{opt} = 1$, and the system performance, with throughput below $1/e$, can be solved from Equation (8b) and Equation (1).

Physically, when p is very small, the send rate B_d is very small (lies in [0,1/N]. See Figure 3). Most times the system is idle and few packets incur collisions. That is, the idle probability P_I is high and the collision probability P_C is low. Increasing p increases B_d and P_C but decreases P_I. Although this leads to more collisions, the number of idle slots also decreases, and the overall system throughput λ_s increases. If with increasing p, the send rate B_d remains below $1/N$ after p reaches 1, the throughput increases monotonically with p (the case $p_{opt} = 1$). Otherwise, the send rate B_d goes beyond $1/N$ after p reaches a certain point ($p_{max} < 1$), the system begins to incur too many collisions, and the throughput begins to drop. That is, in this case, the throughput first increases then decreases, with maximum $1/e$ achieved at p_{max}.

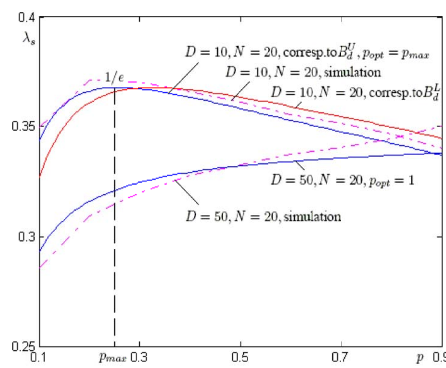

Fig. 4. System Throughput as a Function of Transmission probability

By solving Equation (8b) and (1), Figure 4 shows the numerical results for the throughput λ_s as a function of p when $D = 10$ and $N = 20$ as well as when $D = 50$ and $N = 20$. The first case ($D = 10$) corresponds to the case $p_{opt} < 1$, and as expected, the throughput first increases with p then decreases. The second case ($D = 50$) corresponds to the case $p_{opt} = 1$, and the throughput increases monotonically with p. For comparison, the throughput corresponding to B_d^L is also shown. We can see that the two throughput curves corresponding to B_d^U and B_d^L are very close to each other, which further confirms that using the throughput corresponding to B_d^U to approximate the actual throughput is good.

Figure 4 also plots the simulation results for the original system with the same sets of parameters. It can be seen that the numerical results based on Equation (8b) and (1) match well with the simulation results for the original system.

Moreover, both the numerical results and simulation results for the above cases indicate that the expected collision window decreases monotonically with p, with value below 2 for $p \geq 0.1$. For brevity, the curves are not shown here.

C. Impact of Two-way Propagation Delay on System Performance

Analog to the above analysis for p, it can be shown that B_d^U decreases monotonically with D. For brevity, we omit the details. Furthermore, when D ranges from 1 to infinity, if the solution of condition (9) for D, denoted by D_{max}, lies within $[1, \infty]$, then the ALOHA performance is achieved at D_{max}. In this case both larger D and smaller D lead to a throughout below $1/e$. Otherwise, $D_{max} \notin [1, \infty]$, and the throughput increases monotonically with decreasing D and is always below $1/e$.

Physically, a low throughput can be either due to too many idle slots or too many collisions. Consequently, for normal p and N with $D_{max} \in [1, \infty]$ and starting from D_{max}, increasing D leads to too many idle slots and reducing D leads to too many collisions. Therefore both result in monotonic decreasing of the throughput with D. While for very small p and N with $D_{max} \notin [1, \infty]$, there are always too many idle slots no matter how small D is. Reducing D thus reduces the number of idle slots and always increases the throughput.

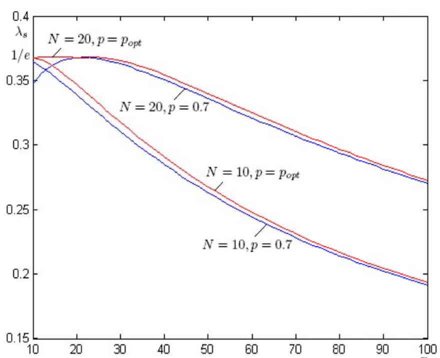

Fig. 5. System Throughput as a Function of Two-way Propagation Delay

Figure 5 shows the numerical results for the throughput λ_s as a function of $D \in [10, 100]$ when $N = 20$ and $p = 0.7$ as well as when $N = 10$ and $p = 0.7$. In the first case ($N = 20$) the ALOHA performance in achieved around $D_{max} \approx 25$, while in the second case ($N = 10$), even when $D = 10$, there are still too many idle slots and the throughput is still below $1/e$.

To illustrate the impact of adjusting protocol parameters on the system performance, for each D in each case, p_{max} and p_{opt} are also calculated from condition (9) and Equation (10) and the resulting throughput is plotted in Figure 5 as well. The figure shows that for large D, $p_{max} \notin [0, 1]$ and $p_{opt} = 1$. The resulting optimal throughput is below $1/e$. While when D is relatively small, the highest possible throughput $1/e$ can always be achieved by setting the transmission probability $p = p_{opt} = p_{max}$. That is, we can always lower the transmission probability p to counterbalance the increasing collisions resulting from smaller D.

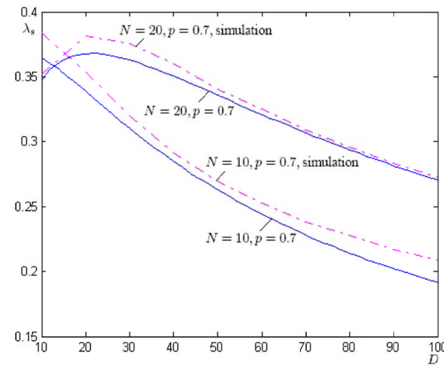

Fig. 6. Comparison between Numerical and Simulation Results - Throughput vs Two-way Propagation Delay

Simulation results for the original system with the same sets of parameters are plotted in Figure 6. For comparison purpose, the numerical results obtained from Equation (8b) and (1) and plotted in Figure 5 are replotted here as well. Again, the figure shows good match between them.

Moreover, the numerical and simulation results for all cases also show that the expected collision window is a nondecreasing function of D and is always below 3. Actually even when D is further increased to 5000, the expected collision window is still below 8. Again for brevity, the curves are not shown here.

D. Impact of Number of Users on System Performance

The analysis for the impact of the number of users N on the system performance is also analog to that for the transmission probability p. Differently, when N increases from 1 to infinity, curve f^U remains the same while curve g moves downwards. Consequently, the intersection of the two curves moves leftwards, and Q^L decreases. Recall that NB_d increases as Q decreases (Equation (2)). Therefore, NB_d increases monotonically with N.

Moreover, it can be easily verified that, for any $p \in [0, 1]$, $D \geq 1$ and $TO \geq \frac{1}{p} + D$ (the actual TO should be at least the average RTT of packets, which is at least $\frac{1}{p} + D$), the solution of condition (9) for N, denoted by N_{max}, is always greater than 1. Therefore, as N ranges from 1 to infinity, NB_d increases from some value below 1 to infinity. Consequently, the system throughput first increases then decreases, with maximum close to $1/e$ achieved at the integer closest to N_{max}.

Physically, larger N makes the packets in the system more "dense", which means less idle slots and more collisions. When N is close to N_{max}, the idle slots and collisions reach a balance and the throughput is close to its highest value $1/e$. Further decreasing N or increasing N results in either too many idle slots, or too many collisions, both of which lower the throughput.

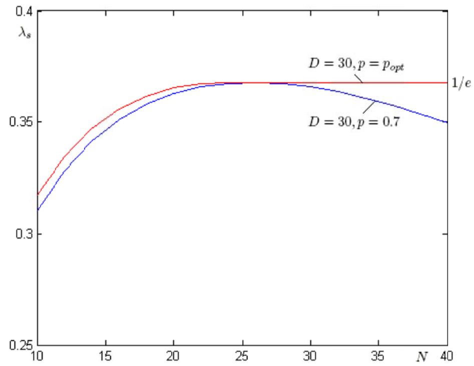

Fig. 7. System Throughput as a Function of the Number of SD Pairs

Figure 7 shows the numerical results for the throughput λ_s as a function of $N \in [10, 40]$ when $D = 30$ and $p = 0.7$ as well as when $D = 30$ and $p = p_{opt}$. As mentioned before, the mean timeout value is set to be $TO = \frac{1}{p} + D + 4\frac{\sqrt{1-p}}{p}$. The throughput of the curve with $p = 0.7$ has maximum value of $1/e$, and first increases then decreases, which confirms the above analysis. The curve with $p = p_{opt}$ further shows that for large enough N, the maximum possible throughput $1/e$ can be achieved by adjusting the transmission probability according to condition (9). That is, in order to achieve higher throughput, we can always lower the transmission probability p to counterbalance the increasing collisions resulting from larger N.

Simulations for the original system with the same sets of parameters are also performed and the results are compared with the numerical results from Equation (8b) and (1) as well. Again, they have a good match. Besides, both of them give the expected collision window size in all cases being below 3. Again for brevity, the curves are not shown here.

In summary, since the numerical curves obtained from the equations match well with the simulation results for the original system, we conclude that the equations can be used to describe the performance of the original system. In particular, given the system and protocol parameters, p, D, N and TO, Equation (1) and (8b) give us the system performance. Condition (9) is a sufficient and necessary condition on the parameters for the system to achieve ALOHA performance, which has maximum possible throughput $1/e$. The optimal transmission probability at which the throughput can achieve its highest value is given in Equation (10). For systems with very small D and/or very large N, the ALOHA performance can always be achieved by setting p to its optimal value. For fixed p, a system with very large D and/or very small N has a smaller throughput than a system with relatively smaller D and/or larger N due to too many idle slots. On the contrary, a system with very large N has a smaller throughput than a system with relatively smaller N due to too many collisions. Furthermore, in all cases in our numerical computations, the expected collision window is below 4. This is because of too many collisions due to random access.

V. Discussions

Our purpose is to analyze the TCP performance over satellite links with random access. We focus on its window flow control mechanism and deliberately disregards other aspects of TCP, such as RTT measurements and estimation and timer granularity. The system analyzed has a window flow control scheme that differs from that in TCP in two aspects: first, we only considered the window evolution in the congestion avoidance phase; and second, we ignored the effect of duplicate ACKs for fast retransmission/recovery. Our simulations show that even with these differences, the analysis gives good prediction on the performance of the original system. Here we discuss the effects of these two differences.

Due to the large number of collisions that result from the random access protocol, the system original with TCP has a very small congestion window size. This is confirmed by the simulations, which shows that it is normally below 4 and below 8 for extremely large D. As a result, the TCP window threshold will also be very small (below four). With such a small window threshold, the window evolution with only congestion avoidance phase is close to that with both slow start and congestion avoidance phase.

Furthermore, the authors in [2] show that for a small congestion window such as eight, fast retransmission/recovery can seldom be entered into if there are multiple losses within the window, which is the case in our system due to collisions. Therefore the window flow control scheme analyzed is a good approximation of that in TCP in our system as well.

For the random access scheme, we did not consider collision recovery at the MAC layer. Over a large delay satellite link, this assumption is also reasonable as any attempt at delayed retransmission will result in a time-out event with very high probability. Hence, it is sensible to leave such retransmissions to higher layers. In addition, we did not take into account the timeout update in TCP. One future direction is therefore to consider collision recovery at the MAC layer as well as to model the timeout update at the transport layer.

References

[1] H. Balakrishnan, V.N. Padmanabhan, S.Seshan and R.H. Katz, “A comparison of mechanisms for improving TCPperformance over wireless links,” IEEE/ACM Transactions on Networking, Vol. 5, pp. 756 -769, December 1997.

[2] K. Fall and S. Floyd, “Simulation-based comparisons of Tahoe, Reno, and SACK TCP,” Computer Communication Review, 1996.

[3] N. Ghani and S. Dixit, “TCP/IP enhancements for satellite networks,” IEEE Communications Magazine, Vol. 37, pp. 64-72, July 1999.

[4] T.V. Lakshman and U. Madhow, “The performance of TCP/IP for networks with high bandwidth-delay products and random loss,” IEEE/ACM Transactions on Networking, Vol. 5, pp. 336 -350, June 1997.

[5] C. Liu and E.H. Modiano, ”On the Interaction of Layered Protocols: The Case of Window Flow Control and ARQ,” Conference on Information Science and System, Princeton, NJ, March, 2002.

[6] M. Meo and C.F. Chiasserini, “Improving TCP over wireless through adaptive link layer setting,” IEEE GLOBECOM ’01, Vol. 3 , pp. 1766 -1770, November 2001.

[7] J. Padhye, V. Firoiu, D.F. Towsley, J.F. Kurose, “Modeling TCP Reno performance: a simple model and its empirical validation,” IEEE/ACM Transactions on Networking, Vol. 8, pp. 133 -145, April 2000.

[8] C. Partridge and T.J. Shepard, “TCP/IP performance over satellite links,” IEEE Network, Vol. 11, pp. 44-49, September/October 1997.