Collaborative SCTP: A Collaborative Approach to Improve the Performance of SCTP over Wired-cum-Wireless Networks

Sha Liu
Department of Computer Science
University of Science and Technology of China
Hefei, Anhui, 230027, China
liusand@mail.ustc.edu.cn

Shoubao Yang
Department of Computer Science
University of Science and Technology of China
Hefei, Anhui, 230027, China
syang@ustc.edu.cn

Weifeng Sun
Department of Computer Science
University of Science and Technology of China
Hefei, Anhui, 230027, China
wfsun@mail.ustc.edu.cn

***Abstract* – Stream Control Transmission Protocol (SCTP) is one of the latest transport layer protocols. It is equipped with many new features, such as multi-stream and multi-homing, and still maintains TCP-like congestion control mechanisms. As TCP, SCTP also fails to differentiate wireless loss from congestion loss, thus its performance over wired-cum-wireless networks suffers from unnecessary congestion window decreases. To improve the performance of SCTP in such a scenario, a new approach with the collaboration of multiple entities and interaction of multiple network layers is proposed in this paper. Simulation experiments conducted through extended ns-2 validated that the proposed approach is adaptive to variable bit error rate (BER) of wireless channel, and can achieve higher goodput along with higher bandwidth utilization efficiency.**

***Keywords* – SCTP, Wireless Networks, 802.11, BER**

I. INTRODUCTION

Stream Control Transmission Protocol (SCTP) is a new transport layer protocol standardized by IETF in RFC 2960[1]. It has almost the same congestion control mechanisms as TCP [2] does. It also interprets all packet losses as congestion losses [3], and then decreases congestion window and packet sending rate. But if the packet sending rate is decreased when the triggering lost packet is actually corrupted, which is very common on wireless channel, the goodput will be impaired and the network resource will be used inefficiently. For this reason, transport layer protocols with TCP-like congestion control mechanisms should be refined to accommodate themselves to such heterogeneous networks.

Till now, many researches have been conducted on the performance of TCP over wired-cum-wireless networks, including I-TCP [4], Snoop [5], TCP-Probing [6], and etc. Since SCTP adopts almost the same congestion control mechanisms as TCP, many schemes originally designed for TCP should be applied to SCTP without many difficulties. But since SCTP has many new features like multi-stream to enhance some applicable applications, there should be some new schemes that utilize these features to enhance the performance of SCTP in such a heterogeneous environment. Unfortunately, research works concerning this aspect have not been published yet. In [7], G. Ye et al. use the existence of Explicit Congestion Notification (ECN) falling in the sending window as the indicator of congestion loss. If the loss is not accompanied by an ECN, the loss is caused by wireless corruption on the first or the last wireless hop. However, this approach does not make any difference between SCTP and TCP, and the packet loss reason indicator can be inaccurate once the ECN is lost due to congestion or corruption. In such cases, the sender may probably continue sending packets at a high rate while actually congestion does occur, resulting in an aggressive TCP-unfriendly behavior.

Contrast against current research status of SCTP over wired-cum-wireless networks, this paper proposed a new approach named Collaborative SCTP to improve the performance of SCTP in this heterogeneous scenario. The approach utilizes the message-orientation and multi-stream features of SCTP, which makes it different from others. It encompasses the collaboration of sender, receiver and the base station or Access Point (AP) in IEEE 802.11 [8], along with the interaction of transport layer and data link layer on wireless stations including base stations.

To validate the improvements brought by such a collaborative approach, we extended the SCTP module in Network Simulator version 2 (ns-2) [9] to incorporate some new functions. The tests are conducted in a typical IEEE 802.11 based wired-cum-wireless network with different parameters. The results of the simulations turned out that this approach is adaptive to variable BER of wireless channel and can achieve higher goodput along with higher bandwidth utilization efficiency. Here the word “efficiency” is defined as the ratio of goodput to all bandwidth consumed by the transmitter no matter it is a wireless station or a base station.

Together with simulation results, the collaborative scheme between multiple entities is presented in Section III, and the whole Collaborative SCTP incorporating interaction between network layers will be discussed in section IV along with the simulation results. Section V then provides the conclusion and some introduction to our future work.

II. ANALYSIS OF SCTP OVER 802.11

A. The Error Recovery of IEEE 802.11

For a unicast frame, IEEE 802.11 requests the receiver to respond with an ACK frame to confirm the correct receipt. If the frame sender fails to receive the ACK frame, it repeats sending the frame until receiving the corresponding ACK frame or until the maximum retransmission count is reached. If all of the attempts to transmit the frame do not result in the receipt of the expected ACK frame, the sender discards the frame. In this way, IEEE 802.11 increases the chance for frames to be received without error. However, it does not guarantee the successful transmission of every frame, so the reliability still needs to be obtained by the retransmission of the lost packets from the SCTP sender.

B. RTS/CTS mechanism of IEEE 802.11

To alleviate frame sending collision, IEEE 802.11 uses the exchange of Request To Send (RTS) and Clear To Send (CTS) to clear the wireless channel. Unfortunately, RTS/CTS mechanism introduces some overhead, so IEEE 802.11 uses the *RTS Threshold* to decide whether to use the RTS/CTS mechanism. If the frame length exceeds *RTS Threshold*, RTS/CTS mechanism is used; otherwise, RTS/CTS exchange is suppressed. Besides that, *RTS Threshold* also determines the maximum transmission attempt for a frame. If the frame length exceeds *RTS Threshold*, *dot11LongRetryLimit* is used to limit the transmission attempt; otherwise, IEEE 802.11 uses *dot11ShortRetryLimit* as the limitation. In ns-2, *RTS Threshold* is set to 0, so *dot11LongRetryLimit*, whose default value is 4, is always used. Setting the *RTS Threshold* to 0 also means that the RTS/CTS mechanism is always used. Using such a mechanism, sending collisions are almost eliminated. It follows that frame sending failures can all be attributed to bit error on the wireless channel. Therefore, the RTS/CTS mechanism will be used to estimate the current BER later.

C. The Impact of Packet Length on the Performance

It is intuitive that the longer the frame, the less possible it can be transmitted without error. In fact, if BER is p, then for a frame with L bits long, the probability P_i that it can be transmitted without error in the ith transmission is

$$P_i = (1-p)^L [1-(1-p)^L]^{i-1} \,. \quad (1)$$

If the maximum transmission count of the specific data link layer is M, the probability P^M of successful frame transmission (briefly success probability) is

$$P^M = \sum_{i=1}^{M} P_i = (1-p)^L \sum_{i=1}^{M} [1-(1-p)^L]^{i-1} \,. \quad (2)$$

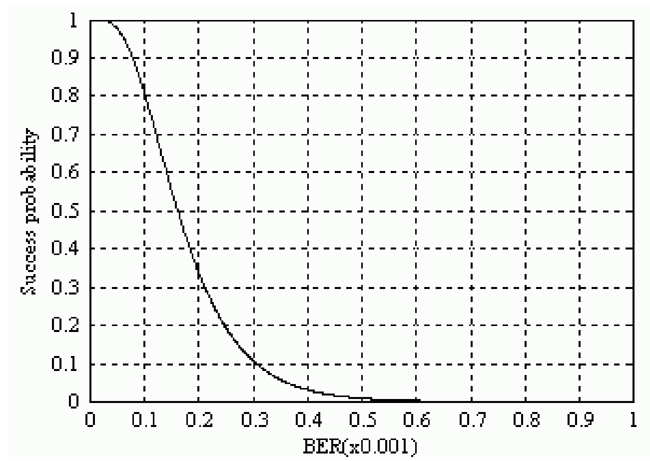

Fig. 1. Success probability as the function of BER

Fig. 2. Simulation topology

For default parameters of IEEE 802.11 in ns-2 and 1500 bytes packet, Fig. 1 depicts the success probability as the function of BER. It can be seen that the success probability drops sharply when BER is higher than 0.00005.

D. Simulation Analysis

In this section, the size of congestion window (briefly cwnd), goodput and bandwidth utilization efficiency (briefly efficiency) are analyzed through simulations. Fig. 2 shows the typical wired-cum-wireless simulation topology, and the simulation parameters are listed in Table I. The wireless station initiates the SCTP association, and retrieves data from the server. The simulation is to get the goodput and efficiency as the functions of discrete BER's from 5×10^{-5} to 2×10^{-4}. For each BER sampled, the simulations were run 10 times to obtain the average performance. The results are shown in Fig. 3-5. Fig. 3 shows the sizes of cwnds as the functions of time for the first runs for BER at 5×10^{-5} and 1×10^{-4}. It can be seen that cwnds oscillate frequently and sharply with small average values, which leads to Fig. 4-5.

TABLE I
SIMULATION PARAMETERS

Parameter	Value
Wired line round trip propagation delay	200ms
Bandwidth of the wired networks	1Mb/s
Wireless MAC protocol	IEEE 802.11
Simulation duration	100 s
SCTP chunk size	260 Bytes
SCTP data packet size	1412 Bytes

Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks (LCN'04)
0742-1303/04 $ 20.00 IEEE

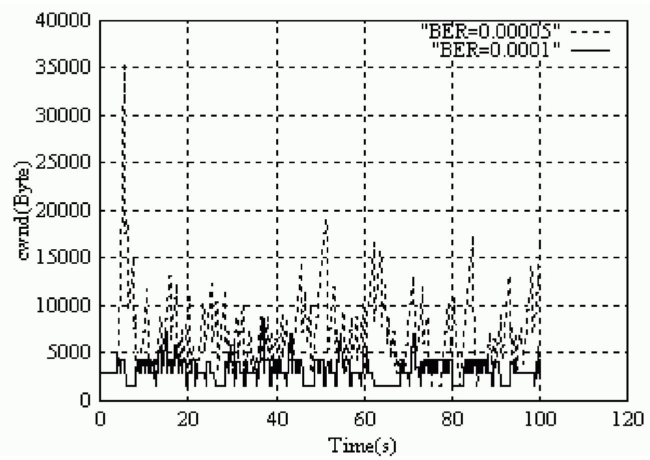

Fig. 3. Cwnds as the function of time for BER 5×10^{-5} and 1×10^{-4}

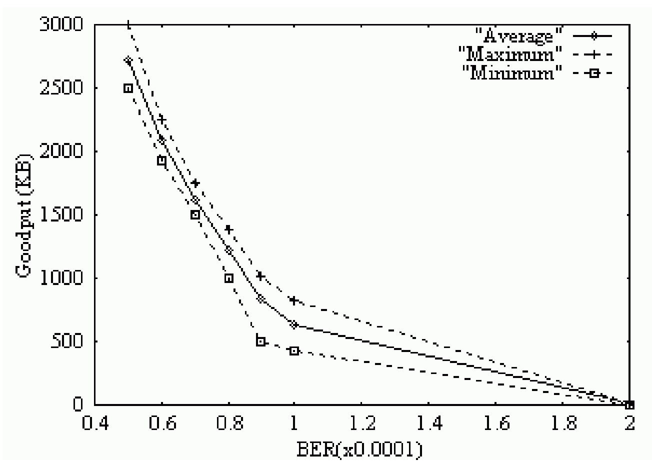

Fig. 4. Goodput as the function of BER

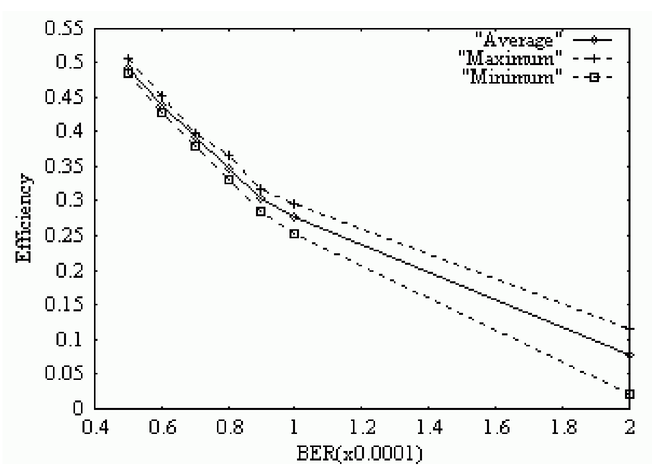

Fig. 5. Efficiency as the function of BER

Type=0	D	N	U	B	E	Chunk Length
Transmission Sequence Number (TSN)						
Stream ID (SID)						Stream Sequence Number (SSN)
Payload Type						
Data						

Fig. 6. Modified data chunk header for disassembly and reassembly

Fig. 4 presents the goodput as the function of BER, and it can be seen that the goodput drops down sharply as BER changes from 1×10^{-5} to 2×10^{-4}. Full records show that the average goodput performance degradations are 76.7% for BER from 5×10^{-5} to 1×10^{-4}, and 99.7% for BER from 5×10^{-5} to 2×10^{-4}. Fig. 5 is about the drop of efficiency, and the degradations are about 0.2 for BER from 5×10^{-5} to 1×10^{-4}, and 0.4 for BER from 5×10^{-5} to 2×10^{-4}. The latter is totally unacceptable. Such performance results motivated the design of Collaborative SCTP presented in Section III and IV.

III. Collaboration of Multiple Entities

A. Disassembly and Reassembly Function

Since the smaller the frame, the greater probability of successful transmission, it is intuitive to send small packets instead of large packets. However, if SCTP sender delivers chunks in small packets, the overhead of IP header and SCTP headers will introduce more overhead for the bit-error-free wired networks. It follows that the proper way is to disassemble a large SCTP packet into small packets at the sending wireless stations including base stations. There are two ways in the literature to achieve this: IP fragmentation and MAC fragmentation. But they have some common limitations. First, they require the successful transmission of all fragments; second, they are both unaware of SCTP. As the description of the next subsection, awareness of SCTP can help sender differentiate wireless loss from congestion loss.

Since SCTP sender can bundle several data chunks in one packet, base station or wireless station can disassemble such a SCTP packet into several small SCTP packets with one or more chunks in each packet. Receiver could interpret each fragmented packet as a complete SCTP packet, but this direct action may trigger too many Selective ACK (SACK) packets destined to the sender. It follows that the receiver should reassemble the fragmented packets into a complete packet. To achieve this goal, the SCTP data chunk header is modified as Fig. 6. D field with 1 bit and N field with 4 bits are added, which denote, respectively, whether it is one of the results from disassembly and how many chunks the original packet contains. If D bit is set to 0, then N field should be set to 0 and ignored by receiver. The disassembly and reassembly algorithms are shown in Fig. 7 and Fig. 8 respectively.

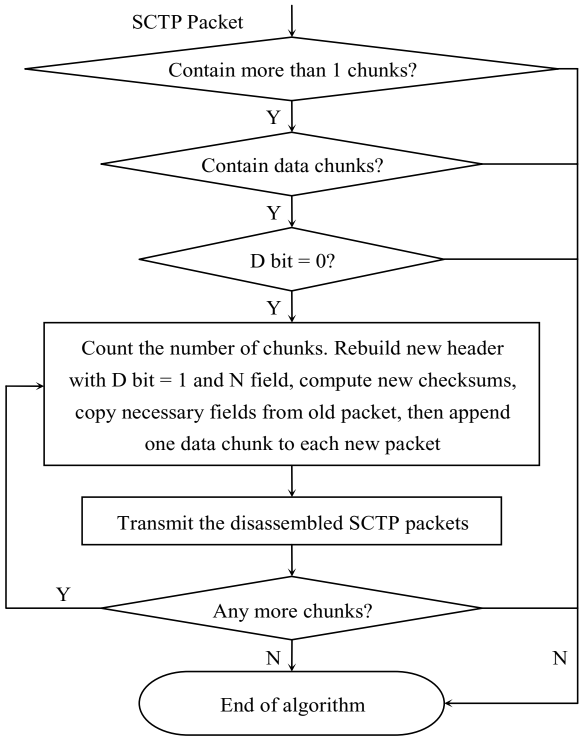

Fig. 7. Disassembly algorithm

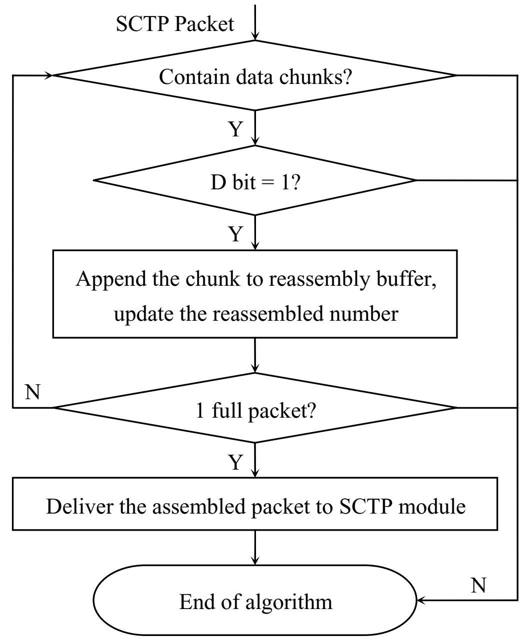

Fig. 8. Reassembly algorithm

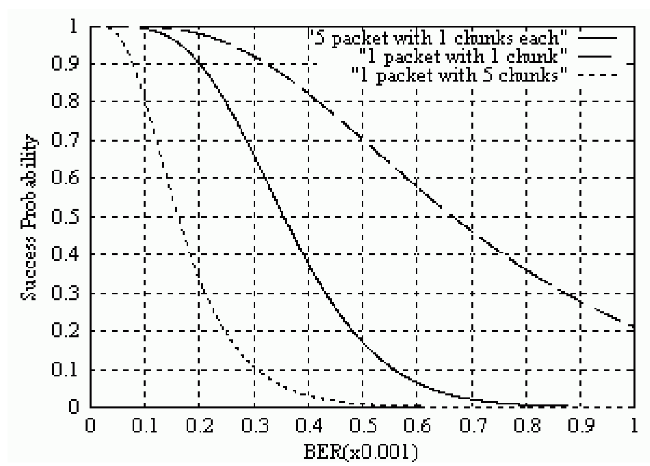

Fig. 9. Probability of successful packet transmission as the function of BER

The disassembly and reassembly mechanisms are designed to increase the probability of successful packet transmission. Fig. 9 shows the success probability of original 1412 bytes packet with 5 chunks, disassembled 308 bytes packet with 1 chunk, and 5 such disassembled packets as a whole. The curves inside Fig. 9 make up of the foundation of SCTP disassembly and reassembly functions.

The disassembly and reassembly functions have special meaning for SCTP transmission. One is that the chunk in each disassembled packet may fall into different logical streams, so the disassembly and reassembly alleviated the Head of Line Blocking (HLB) in a finer granularity. Second, it is useful for sender to differentiate the wireless loss from congestion loss, as described in the next subsection.

B. Log of Chunks Bundle and Loss Differentiation

The disassembly and reassembly functions can greatly increase the transmission success probability of SCTP packets, but they can not eliminate all wireless losses. Other makeup mechanisms are still needed at the sender to differentiate the reason of packet loss. To design such mechanisms, first notice that, although some packets originating from disassembling a large packet are possible to get corrupted, the chance is very small that all these packets get damaged.

For example, packet A contains chunks with TSN 1-5, and packet B contains chunks 6-10. If only the disassembled packet containing chunk 3 is lost, the receiver will issue SACK packet to acknowledge the receipt of chunks 1-2 and chunks 4-10. If the sender can recall that chunks 1-5 are originally bundled in a SCTP packet, it will be sure that the original packet has been undergone disassembly, and the loss of chunk 3 is caused by the damage on the wireless channel, not by congestion. This discussion motivated the idea that the sender logs the chunks bundles and uses the records to distinguish between these two kinds of losses.

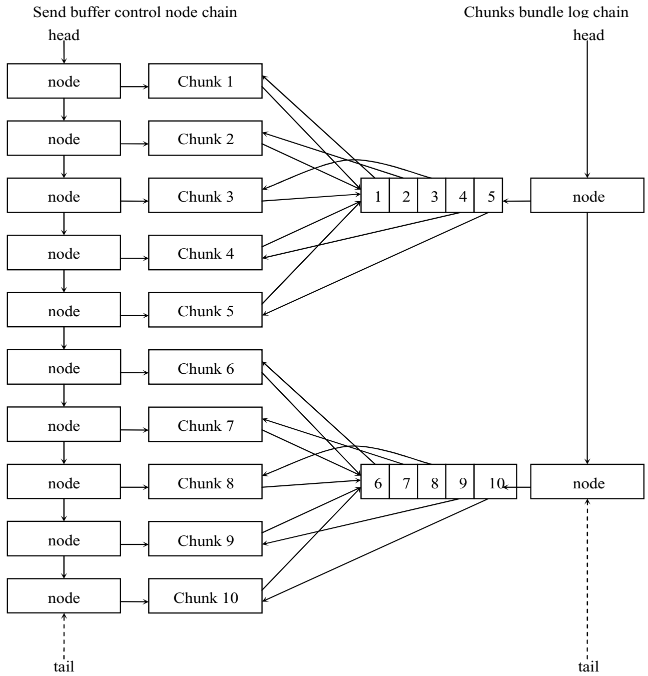

Fig. 10. Structure of chunks bundle log management

SCTP module in ns-2 uses a send buffer control node to manage each chunk that has been sent but has not been acknowledged. Each control node contains the necessary information about the chunk including the pointer to the chunk buffer, retransmission mark, flag about the eligibility of fast retransmission, and etc. A chain is used to link all these control nodes together. The chunks bundle log is designed based on a similar structure. Each SCTP packet sent out is managed by a similar node which points to an array. Each element in the array manages a chunk by maintaining its TSN and the pointer to its buffer. At the same time, a pointer to the chunks bundle log node is added to each original send buffer control node, as Fig. 10 exhibits.

After receiving SACK packets, checks should be done to ensure the chunks managed by such a node have all been acknowledged. If it is the case, then the node is freed. If all chunks in a chunks bundle log node are reported to have been lost, it is assumed a congestion loss, which may be wrong but is conservative enough to be TCP-friendly. Otherwise, the sender is sure that the chunk loss is a wireless loss, even without duplicate ACK's in fast retransmission. So the sender marks the lost chunk for retransmission, and retransmits it without cwnd reduction. As chunk buffers may have been released before their TSN's has gone through such checks, it is necessary to maintain the TSN's in array elements while not just looking into the chunk buffers.

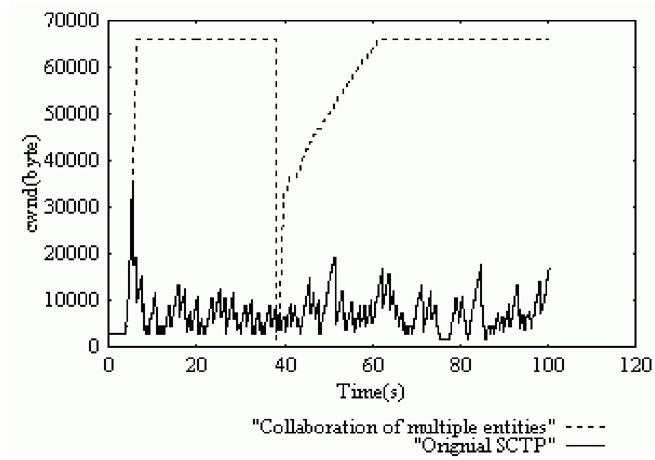

Fig. 11. Congestion window size when BER=0.00005

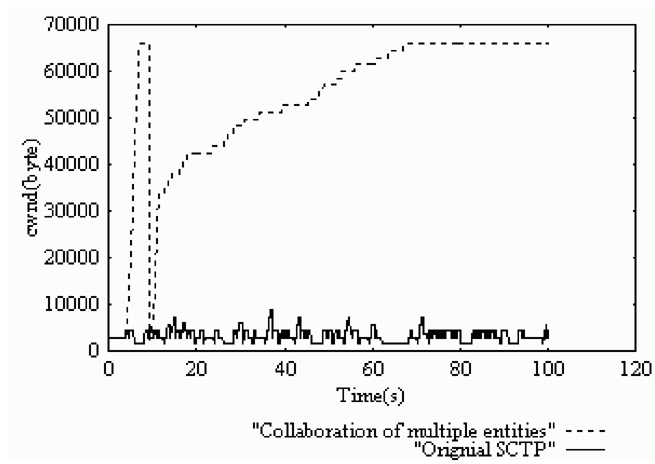

Fig. 12. Congestion window size when BER=0.0001

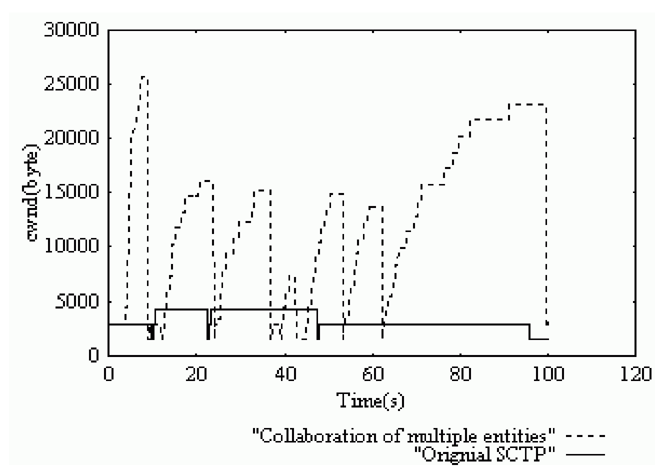

Fig. 13. Congestion window size when BER=0.0002

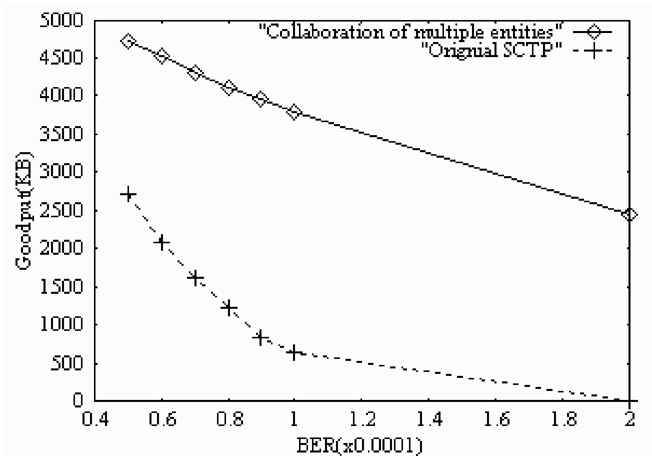

Fig. 14. Goodput as the function of BER

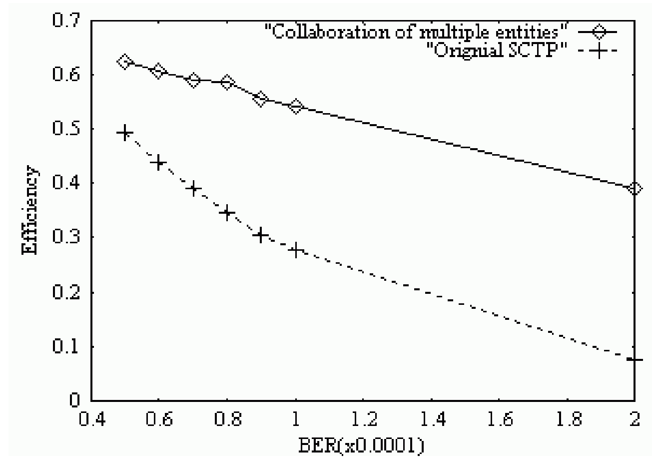

Fig. 15. Efficiency as the function of BER

C. Simulation Results

Using collaboration of Subsection A and B, congestion window, goodput and bandwidth utilization efficiency are also analyzed through simulations. The topology is the same as Fig. 2, and parameters are listed in Table I.

Fig. 11~13 present the variations of congestion window in the simulation duration when BER is 5×10^{-5}, 1×10^{-4} and 2×10^{-4} respectively. It can be seen that the oscillation of congestion window is much steadier if the collaboration of multiple entities is utilized, and congestion window can be expanded to a much larger value. Hence the improvement of goodput and bandwidth utilization efficiency can be anticipated. Actually, Fig. 14 and Fig. 15 plot the goodput and efficiency as functions of BER. It is clear that the goodput and efficiency are all improved greatly with the collaboration scheme. In fact, the higher the BER, the more performance improvement can be achieved. The goodput improvement arrives at 73.4% when BER is 5×10^{-5}, and 428.7% at BER of 1×10^{-4}. Efficiency is increased by more than 20% at BER of 5×10^{-5}, about 50% at BER of 1×10^{-4}, and more than 300% at BER of 2×10^{-4}. As BER increases, the drops are quite slow, showing a much more robustness against high BER.

D. Some Discussions

Disassembly and reassembly at base station are important for the system discussed above, but they somewhat violates the end-to-end argument. However, the performance gains are so significant that it is hard to reject the idea. In addition, R. Haas et al. presented the idea of moving network protocol analyses from end host to networks using network processors in [10], and took the analysis of SCTP as an example. This idea consolidated ours: equipping base stations with these functions through network processors is actually preferable.

The second concern is about the security. Base station may not be able to decrypt the packets with IPSec ESP protection, but sender can encrypt messages at transport layer or application layer using SSL or SSH. Security per chunk contents sounds even more reasonable as the independence between difference streams in the same SCTP association.

IV. INTERACTION OF MULTIPLE LAYERS

SCTP disassembly and reassembly functions together with the packet loss differentiation mechanism of the sender make up the collaboration of multiple entities. Simulations show that the goodput and bandwidth utilization efficiency can be greatly improved especially at relatively high BER. But if the BER is so low that the wireless losses are only rare events, the disassembly function can bring in more IP/SCTP protocol overhead, more RTS/CTS overhead and more ACK frame overhead, which may result in lower performance. For this reason, there should be some approach to monitor the wireless channel and estimate current BER. Furthermore, the disassembly mechanism is activated when the measured BER is high, and deactivated when estimator indicates a low BER.

A. Estimation of Current BER

As early statements, the exchange of RTS/CTS can clear the wireless channel during the atomic transmission. In this way, the corrupt of frames can be attributed to the wireless random error with high probability. Equation (2) states the relationship between BER and transmission count. The data link layer can continuously record the transmission count of each frame sent, and then use a certain number of latest samples to estimate current BER. With continuous n latest independent samples that are all L bits long, the transmission counts $s=(s_1, s_2, \dots, s_n)$ conform to the sample distribution in equation (3).

Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks (LCN'04)
0742-1303/04 $ 20.00 IEEE

$$P(s)=\prod_{i=1}^{n}(1-p)^{L}[1-(1-p)^{L}]^{s_i-1}\ . \quad (3)$$

View the sample distribution as the function of independent variable *p*, denoted by *f*(*p*; *s*), then Equation (3) is a likelihood function, and its logarithm counterpart is

$$\log f(p;s)=\sum_{i=1}^{n}\{L\log(1-p)+(s_i-1)\log[1-(1-p)^{L}]\}, s_i=1,2,...,M+1, \quad (4)$$

where the maximum value of s is $M+1$ because the system logs the transmission count as $M+1$ when maximum transmission count is reached. Logarithm likelihood function (4) leads to the BER Maximum Likelihood Estimator (MLE)

$$\hat{p}=1-\bar{s}^{-\frac{1}{L}}, \quad (5)$$

where $\bar{s}$ is the mean of all samples.

But it is not always the case that all of sample frames are of the same length. To simplify the solution of MLE, mean length $\bar{L}$ of sample frames is used instead of *L*. It follows that the MLE of BER is

$$\hat{p}=1-\bar{s}^{-\frac{1}{\bar{L}}}. \quad (6)$$

B. Interaction of Multiple Layers

A well known error model for wireless channel is represented by a two-state Markov chain. One state means relatively high BER, and the other one stands for relatively low BER. The time that wireless channel lingers in one state conforms to an exponential distribution. At the end of staying in one state, the system has the probability P_{ij} to transit to another state, where *i* represents current state and $i \neq j$.

The two-state Markov chain model can be used to represent states with the high BER and low BER. If current BER estimated using approximate MLE is higher than a critical value, then the disassembly function is activated. If current BER estimated is lower than another lower critical value, it is deactivated. Since the distribution of the time of staying in one state is exponential distribution, which is memoryless, the presumption that the wireless channel will continue lingering in current estimated state is reasonable. To avoid too frequent activation and deactivation, the critical interval bounded by the two values should be carefully chosen. The criterion is that the performance gain or loss is minor no matter the disassembly function is enabled or is disabled when BER falls into the critical interval. After numerous simulations with different chunk sizes, (3×10^{-5}, 4×10^{-5}) is adopted as the critical interval.

The estimation is conducted by MAC module, but the disassembly function is done by the SCTP module. Therefore, the activation and reverse involve the interaction of data link layer and transport layer. This interaction and collaboration of multiple entities (Section III) compose the whole Collaborative SCTP. In this way, the high probability of success transmission from disassembly and the low overhead from original SCTP are both used in proper environments.

C. Simulation Results

The performances are also analyzed via simulation. The simulation scenario is also as Fig. 2 and Table I illustrate. However, the BER is not consistent here; instead, a two-state Markov chain is used to emulate the real environment. Here error models are represented by parameters of BER pairs, mean time of staying in each state, and the state transition probability matrix. BER pairs include (1×10^{-5}, 2×10^{-5}), (5×10^{-5}, 6×10^{-5}), (1×10^{-5}, 1×10^{-4}) and (2×10^{-5}, 8×10^{-5}). The first parameter pair means a usually good channel, the second pair represents a usually bad channel, and the last two denote the environment with large BER variation. The total simulation duration is 100s, but the mean time of staying in each state is 10s. In addition, the state transition probabilities are all 0.5. To accurately estimate BER, large sample size is necessary. But to reflect the real environment in time, the sample size can not be large. In this paper, the sample size is fixed to 32 frames for simplicity and tradeoff.

All simulation results are summarized in Table II and Table III. It can be seen that, when BER pair falls in one side out of the critical interval, the performances of Collaborative SCTP are slightly lower than those that simply use or do not use disassembly function. Such degradation can be attributed to some error estimations. However, for BER pairs outside of both sides the critical interval, the goodput gain is much more significant when using Collaborative SCTP, and efficiency is kept very close to the highest one of using or not using disassembly. Recall that the higher the BER, the more performance gain can be obtained. It can be stated that it is beneficial to deploy Collaborative SCTP especially in fields with high BER.

D. Some Discussions

The first concern about the estimator of BER is that its calculation includes some expensive operations including division and power. Therefore, we suggest the using of a two dimension table containing the BER estimator for some common mean frame length values and some common mean transmission count values. Then the estimator calculation is simplified as a table search with interpolation if necessary.

In addition, as different areas within the coverage of a certain AP may be in different BER states, estimating BER for each area or station may be needed.

The property of Collaborative SCTP to be TCP-friendly is also worth discussing. When BER is high, TCP and original SCTP can not use up their fair bandwidth shares of the base station. Therefore, such moderate overhead is allowed. When BER is low, the disassembly function will be turn off in time, and the fairness infection is only slight.

TABLE II
SIMULATION RESULTS OF GOODPUT (BYTE) UNDER DIFFERENT BER PAIRS

BER	1×10^{-5}, 2×10^{-5}	5×10^{-5}, 6×10^{-5}	1×10^{-5}, 1×10^{-4}	2×10^{-5}, 8×10^{-5}
Original SCTP	7386314	2393924	3436940	3615950
Collaboration of multiple entities	5371132	4599062	4567888	4666168
Multiple layers interaction added	7330856	4343586	5582694	5272826
Improvement	-0.7%	-5.6%	22.2%	13.0%

TABLE III
SIMULATION RESULTS OF EFFICIENCY UNDER DIFFERENT BER PAIRS

BER	1×10^{-5}, 2×10^{-5}	5×10^{-5}, 6×10^{-5}	1×10^{-5}, 1×10^{-4}	2×10^{-5}, 8×10^{-5}
Original SCTP	74.5%	46.5%	61.8%	58.1%
Collaboration of multiple entities	69.1%	61.5%	62.0%	62.7%
Multiple layers interaction added	74.3%	58.1%	65.1%	62.5%
Improvement	-0.2%	-2.6%	3.1%	1.9%

V. CONCLUSIONS AND FUTURE WORK

This paper presented the Collaborative SCTP comprising of SCTP-aware disassembly and reassembly functions, log of chunks bundle, mechanism to differentiate wireless loss from congestion loss, wireless channel monitor, and the system dynamic configuration. Simulation results show that such a collaborative approach can greatly improve the goodput and bandwidth utilization efficiency, especially in high BER environment. Besides the direct performance gain illustrated in Section III and IV, this approach is unique to SCTP since it utilizes the message-orientation and multi-stream features, and is useful for alleviating the HLB in a finer granularity than original SCTP. Henceforth, it supercedes the functions of IP fragmentation and MAC fragmentation for SCTP optimization in wired-cum-wireless networks.

It should also be aware that the BER estimation is quite coarse, which is caused by the following facts:

a. To simplify the calculation of BER MLE, the mean length of all sample frames is used as an approximation;

b. For frames that ultimately failed to be transmitted and discarded by the data link layer, the system records its transmission count as the maximum transmission count plus one, which is an approximation;

c. Relatively small sample size;

d. Chances are that some moving wireless stations are unaware of certain RTS/CTS exchanges, which may cause sending collisions and overestimation of BER.

Our future work is first to refine the BER estimator aiming at the four points listed above. In addition, the incorporation of Collaborative SCTP with some TCP enhancing schemes in the literature like Snoop, TCP-Probing and ECN should be worthwhile. Besides that, performance studies of such a collaborative approach in mobile ad hoc networks or other wireless access technologies are also interesting and necessary. We hope to present these results in the future.

REFERENCES

[1] R. Stewart et al., "Stream Control Transmission Protocol", RFC 2960, Oct. 2000.

[2] ISI/USC, "Transmission Control Protocol", RFC 793, Sept. 1981.

[3] V. Jacobson and M. Karels, "Congestion Avoidance and Control", Proceedings of SIGCOMM 1988, pp. 314-329, Aug. 1988.

[4] A. Bakre and B.R. Badrinath, "I-TCP: Indirect TCP for Mobile Hosts", Proceedings of ICDCS 1995, pp. 136-143, May 1995.

[5] H. Balakrishnan, S. Seshan, and R.H. Katz, "Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks", ACM Wireless Networks Journal, pp. 469-481, December 1995.

[6] V. Tsaoussidis and H. Badr, "TCP-Probing: Towards an Error Control Schema with Energy and Throughput Performance Gains", Proceedings of ICNP 2000, Pages 12-21, Nov. 2000.

[7] G. Ye, T. Saadawi, and M. Lee, "Improving Stream Control Transmission Protocol Performance over Lossy Links", IEEE Journal on Selected Areas in Communications, Volume 22, Issue 4, pp. 727-736, May 2004.

[8] IEEE 802.11 WG, "ANSI/IEEE Std 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", 1999.

[9] The Network Simulator: ns-2, http://www.isi.edu/nsnam/ns/.

[10] R. Haas et al., "Creating Advanced Functions on Network Processors: Experience and Perspectives", IEEE Network, Volume 17, Issue 4, pp. 46-54, July-Aug. 2003.