

# Concurrent Multipath Transfer Using SCTP Multihoming Over Independent End-to-End Paths

Janardhan R. Iyengar, Paul D. Amer, and Randall Stewart

**Abstract**—Concurrent Multipath Transfer (CMT) uses the Stream Control Transmission Protocol’s (SCTP) multihoming feature to distribute data across multiple end-to-end paths in a multihomed SCTP association. We identify three negative side-effects of reordering introduced by CMT that must be managed before full performance gains of parallel transfer can be achieved: (i) unnecessary fast retransmissions by a sender, (ii) overly conservative cwnd growth at a sender, and (iii) increased ack traffic due to fewer delayed acks by a receiver. We propose three algorithms which augment and/or modify current SCTP to counter these side-effects. Presented with several choices as to where a sender should direct retransmissions of lost data, we propose five retransmission policies for CMT. We demonstrate spurious retransmissions in CMT with all five policies, and propose changes to CMT to allow the different policies. CMT is evaluated against AppStripe, an idealized application that stripes data over multiple paths using multiple SCTP associations. The different CMT retransmission policies are then evaluated with varied constrained receive buffer sizes. In this foundation work, we operate under the strong assumption that the bottleneck queues on the end-to-end paths used in CMT are independent.

**Index Terms**—Load balancing, load sharing, multipath, SCTP, transport layer, end-to-end.

## I. INTRODUCTION

A host is multihomed if it can be addressed by multiple IP addresses, as is the case when the host has multiple network interfaces. Though feasibility alone does not determine adoption of an idea, multihoming is increasingly economically feasible and can be expected to be the rule rather than the exception in the near future, particularly when fault tolerance is crucial. Multihomed nodes may be simultaneously connected through multiple access technologies, and even multiple end-to-end paths to increase resilience to path failure. For instance, a mobile user could have simultaneous Internet connectivity via a wireless local area network using 802.11b and a wireless wide area network using GPRS.

We propose using *Concurrent Multipath Transfer (CMT)* between multihomed source and destination hosts to increase an application’s throughput. CMT is the concurrent transfer of

new data from a source to a destination host via two or more end-to-end paths.

The current transport protocol workhorses, TCP and UDP, do not support multihoming; TCP allows binding to only one network address at each end of a connection. At the time TCP was designed, network interfaces were expensive components, and hence multihoming was beyond the ken of research. Changing economics and an increased emphasis on end-to-end fault tolerance have brought multihoming within the purview of the transport layer. While concurrency can be arranged at other layers (as discussed in Sections IV-D and VI), the transport layer has the best knowledge to estimate end-to-end paths’ characteristics.

The Stream Control Transmission Protocol (SCTP) [1], [2] is an IETF standards track protocol that natively supports multihoming at the transport layer. SCTP multihoming allows binding of one transport layer *association* (SCTP’s term for a connection) to multiple IP addresses at each end of the association. This binding allows a sender to transmit data to a multihomed receiver through different destination addresses. Simultaneous transfer of new data to multiple destination addresses is currently not allowed in SCTP due primarily to insufficient research. This research attempts to provide that needed work.

Though CMT uses SCTP in our analysis, our goal is to study CMT at the transport layer in general. The issues and algorithms considered in this research would apply to any multihome-aware transport protocol. We chose SCTP primarily due to lack of mature multihoming mechanisms in any other practical transport layer protocol<sup>1</sup>, and partly due to our expertise with it.

Following preliminary concepts and terminology in Section II, Section III specifies three algorithms resulting in CMT<sub>scd</sub> - a protocol that uses SCTP’s multihoming feature for *correctly* transferring data between multihomed end hosts using multiple independent end-to-end paths. A CMT sender can direct retransmissions to one of several destinations that are receiving new transmissions. In Section IV, we present an evaluation of CMT versus an “idealized” hypothetical application which stripes data across multiple SCTP associations (AppStripe). We also propose and evaluate five retransmission policies for CMT. We conclude our discussion of CMT in Section V, and discuss related work in Section VI.

<sup>1</sup>The Datagram Congestion Control Protocol (DCCP) [3] does provide “primitive multihoming” at the transport layer, but only for mobility support. DCCP multihoming is useful only for connection migration, and cannot be leveraged for CMT.

J. Iyengar and P. Amer are with the Protocol Engineering Lab, CIS Dept., University of Delaware, Newark, DE 19716 USA (email: iyengar@cis.udel.edu; amer@cis.udel.edu).

R. Stewart is with the Internet Technologies Division at Cisco Systems (email: rrs@cisco.com).

Prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

Supported in part by University Research Program of Cisco Systems, Inc.

## II. PRELIMINARIES

We first overview several ideas and mechanisms used by SCTP relevant to this research; some are compared with TCP to highlight similarities and differences. SCTP is defined in RFC2960 [2] with changes and additions included in the Implementer’s Guide [1]. An SCTP packet, or protocol data unit (PDU), consists of one or more concatenated building blocks called *chunks*: either control or data. For the purposes of reliability and congestion control, each data chunk in an association is assigned a unique Transmission Sequence Number (TSN). Since SCTP is message-oriented and chunks are atomic, TSNs are associated with chunks of data, as opposed to TCP which associates a sequence number with each data octet in the bytestream. In our simulations, we assume one data chunk per PDU for ease of illustration; each PDU thus carries, and is associated with a single TSN.

SCTP uses a selective ack scheme similar to SACK TCP [4]. SCTP’s congestion control algorithms are based on RFC 2581 [5], and include SACK-based mechanisms for better performance. Similar to TCP, SCTP uses three control variables: a receiver’s advertised window (rwnd), a sender’s congestion window (cwnd), and a sender’s slow start threshold (ssthresh). However, unlike TCP’s cwnd which reflects which data to send, SCTP’s cwnd dictates only how much data can be sent. SCTP’s cwnd is similar in function to SACK TCP’s *pipe* variable [6]. In SCTP, rwnd is shared across an association. Since an SCTP association allows multihomed source and destination endpoints, a source maintains several parameters on a *per destination* basis: cwnd, ssthresh, and roundtrip time (RTT) estimates. An SCTP sender also maintains a separate retransmission timer per destination. RFC 2960 does not allow a sender to simultaneously send *new* data on multiple paths. New data must be sent to a single *primary destination*, while retransmissions may be sent to any alternate destination.

In this work, we operate under the assumption that the bottleneck queues on the end-to-end paths used in CMT are independent of each other. Overlap in the paths is acceptable, but again bottlenecks are assumed independent. We recognize that this strong assumption certainly will not hold for all multihomed associations. Continued work [7] will investigate CMT without this assumption using end-to-end shared bottleneck detection techniques [8]–[11].

A note on language and terminology. A reference to “cwnd for destination X” means the cwnd maintained at the sender for destination X, and “timeout on destination X” refers to the expiration of a sender’s retransmission timer maintained for destination X. Since bottleneck queues on the end-to-end paths are assumed independent, each destination in our topology uniquely maps to an independent path. Thus, “cwnd for destination X” may be used interchangeably with “cwnd for path Y”, where path Y ends in destination X. SCTP acks carry cumulative and selective ack (also called *gap ack*) information and are called SACKs. In the paper, sometimes “SACK” is used rather than “ack” to emphasize when an ack carries both cumulative and selective acks.

The simulations presented in this paper use the University of Delaware’s SCTP module for ns-2 [12], [13].

## III. CMT ALGORITHMS

As is the case with TCP [14]–[17], reordering introduced in an SCTP flow degrades performance. When multiple paths being used for CMT have different delay and/or bandwidth characteristics, significant packet reordering can be introduced in the flow by a CMT sender. Reordering is a natural consequence of CMT, and is difficult to eliminate in an environment where the end-to-end path characteristics are changing or unknown a priori, as in the Internet. In this section, we identify and resolve the negative side-effects of sender-introduced reordering by CMT in SCTP.

To demonstrate the effects of reordering introduced in SCTP by CMT, we use a simple simulation setup. Two dualhomed hosts, sender *A* with local addresses  $A_1, A_2$ , and receiver *B* with local addresses  $B_1, B_2$ , are connected by two separate paths: Path 1 ( $A_1 - B_1$ ), and Path 2 ( $A_2 - B_2$ ) having end-to-end available bandwidths 0.2 Mbps and 1 Mbps, respectively. The roundtrip propagation delay on both paths is 90 ms, roughly reflecting the U. S. coast-to-coast delay. CMT sender *A* sends data to destinations  $B_1$  and  $B_2$  concurrently, and uses a scheduling algorithm that sends new data to a destination when allowed by the corresponding cwnd.

The simulation results described in this section (Figures 1 and 5) both show cwnd evolution with time. The figures have four curves, which show the CMT sender’s (1) observed cwnd evolution for destination  $B_1$  (+), (2) observed cwnd evolution for destination  $B_2$  (×), (3) calculated aggregate cwnd evolution (sum of (1) and (2)) ( $\Delta$ ), and (4) expected aggregate cwnd evolution (–). This last curve represents our initial performance goal for CMT - the sum of the cwnd evolution curves of two independent SCTP runs, using  $B_1$  and  $B_2$  as the primary destination, respectively.




Fig. 1. CMT with SCTP: Evolution of the different cwnds

Figure 1 shows how, when using SCTP without any modifications, CMT reordering significantly hinders both  $B_1$  and  $B_2$ ’s cwnd growth. (Normally cwnd reductions are seen when a sender detects loss, but for Figure 1, no packet loss was simulated.) The aggregate cwnd evolution ( $\Delta$ ) is significantly below the expected aggregate cwnd evolution (–).

To explain our algorithms which enable correct CMT using SCTP, we first introduce some notation. CMT refers to a host performing concurrent multipath transfer using current SCTP.  $CMT_s$ ,  $CMT_c$ , and  $CMT_d$  refer to a host performing CMT with the Split Fast Retransmit (SFR) algorithm (Section III-A), the Cwnd Update for CMT (CUC) algorithm (Section III-B) and the Delayed Ack for CMT (DAC) algorithm (Section III-C), respectively. Using more than one subscript means inclusion of more than one algorithm. For instance,  $CMT_{sc}$  refers to a host performing CMT with the SFR and CUC algorithms.

We identify and resolve three negative side-effects of reordering introduced by CMT that must be managed before the full performance gains of parallel transfer can be achieved: (i) unnecessary fast retransmissions at the sender (Section III-A), (ii) reduced cwnd growth due to fewer cwnd updates at the sender (Section III-B), and (iii) increased ack traffic due to fewer delayed acks (Section III-C) [18].

#### A. Preventing Unnecessary Fast Retransmissions

When reordering is observed, a receiver sends gap reports (i.e., gap acks) to the sender which uses the reports to detect loss through a fast retransmission procedure similar to the one used by TCP [2], [5]. With CMT, unnecessary fast retransmissions can be caused due to reordering [19], with two negative consequences: (1) since each retransmission is assumed to occur due to a congestion loss, the sender reduces its cwnd for the destination on which the retransmitted data was outstanding, and (2) a cwnd overgrowth problem explained in [20] causes a sender’s cwnd to grow aggressively for the destination on which the retransmissions are sent, due to acks received for original transmissions. In Figure 1, each cwnd reduction observed for  $B_1$  and  $B_2$  is due to an unnecessary fast retransmission by the sender. These unnecessary retransmissions seen by CMT are due to sender-introduced reordering, and not spurious retransmissions due to network effects [21], [22].

Conventional interpretation of a SACK chunk in SCTP (or ack with SACK option in TCP) is that gap reports imply possible loss due to congestion. The probability that a TSN is lost, as opposed to being reordered, increases with the number of gap reports received for that TSN. Due to reordering, a CMT sender needs additional information to infer loss. Gap reports *alone* do not (necessarily) imply loss; but a sender can infer loss using gap reports *and* knowledge of each TSN’s destination.

**Algorithm Details:** The proposed solution to address the side-effect of incorrect cwnd evolution due to unnecessary fast retransmissions is the Split Fast Retransmit (SFR) algorithm, shown in Figure 2. This algorithm extends a previous incarnation which could not handle *cycling changeover* [19]. SFR introduces a *virtual queue* per destination within the sender’s retransmission queue. A sender then deduces missing reports for a TSN correctly using SACK information in conjunction with state maintained about the transmission destination for each TSN in the retransmission queue. Thus, SFR enables a multihomed sender to correctly apply the fast retransmission procedure on a per destination basis. An advantage of SFR is that only the sender’s behavior is affected.

The SFR algorithm introduces two additional variables per destination at a sender:

- 1) *highest\_in\_sack\_for\_dest* - stores the highest TSN acked per destination by the SACK being processed.
- 2) *saw\_newack* - a flag used during the processing of a SACK to infer the causative TSN(s)’s destination(s). Causative TSNs for a SACK are those TSNs which caused the SACK to be sent (or TSNs that are being acked in this SACK for the first time).

In Figure 2, step (2) infers the destinations that the newly acknowledged TSNs were sent to, and sets *saw\_newack* to TRUE for those destinations. Step (3) tracks on a per destination basis, the highest TSN that was being acked. Step (4) uses information gathered in steps (2) and (3) to aid in inferring missing TSNs. Two conditions in step (4) ensure correct missing reports: (a) TSNs to be marked should be outstanding on the same destination(s) as TSNs which have been newly acked, and (b) at least one TSN, sent later than the missing TSN, should have reached the receiver *on the same destination address*.

#### B. Avoiding Reduction in Cwnd Updates

The cwnd evolution algorithm for SCTP [2] (and analogously for SACK TCP [4], [5]) allows growth in cwnd only when a new cum ack is received by a sender. When SACKs with unchanged cum acks are generated (say due to reordering) and later arrive at a sender, the sender does not modify its cwnd. This mechanism again reflects the conventional view that a SACK which does not advance the cum ack indicates possibility of loss due to congestion.

Since a CMT receiver naturally observes reordering, many SACKs are sent containing new gap reports but not new cum acks. When these gaps are later acked by a new cum ack, cwnd growth occurs, but only for the data newly acked in the most recent SACK. Data previously acked through gap reports will not contribute to cwnd growth. This behavior prevents sudden increases in the cwnd resulting in bursts of data being sent. Even though data may have reached the receiver “in-order per destination”, without changing the current handling of cwnd, the updated cwnd will not reflect this fact.

This inefficiency can be attributed to the current design principle that the cum ack in a SACK, which tracks the latest TSN received in-order at the receiver, applies to an entire association, not per destination. TCP and current SCTP use only one destination address at any given time to transmit new data to, and hence, this design principle works fine. Since CMT uses multiple destinations simultaneously, cwnd growth in CMT demands tracking the latest TSN received in-order *per destination*, information not coded directly in a SACK.

We propose a cwnd growth algorithm to track the earliest outstanding TSN *per destination* and update the cwnd, even in the absence of new cum acks. The algorithm uses SACKs and knowledge of transmission destination for each TSN to deduce in-order delivery per destination. The crux of the CUC algorithm is to track the earliest outstanding data *per destination*, and use SACKs which ack this data to update the corresponding cwnd. In understanding our proposed solution,

```

On receipt of a SACK containing gap reports [Sender side behavior]:
1)  $\forall$  destination addresses  $d_i$  initialize  $d_i.saw\_newack = \text{FALSE}$ ;
2) for each TSN  $t_a$  being acked that has not been acked in any SACK thus far do
   let  $d_a$  be the destination to which  $t_a$  was sent;
   set  $d_a.saw\_newack = \text{TRUE}$ ;
3)  $\forall$  destinations  $d_n$ , set  $d_n.highest\_in\_sack\_for\_dest$  to highest TSN being newly acked on  $d_n$ ;
4) to determine whether missing report count for a TSN  $t_m$  should be incremented:
   let  $d_m$  be the destination to which  $t_m$  was sent;
   if ( $d_m.saw\_newack = \text{TRUE}$ ) and ( $d_m.highest\_in\_sack\_for\_dest > t_m$ ) then
     increment missing report count for  $t_m$ ;
   else do not increment missing report count for  $t_m$ ;

```

Fig. 2. SFR Algorithm – Eliminating unnecessary fast retransmissions

```

At beginning of an association [Sender side behavior]:
 $\forall$  destinations  $d$ , reset
   $d.find\_pseudo\_cumack = \text{TRUE}$ ;
On receipt of a SACK [Sender side behavior]:
1)  $\forall$  destinations  $d$ , reset  $d.new\_pseudo\_cumack = \text{FALSE}$ ;
2) if the SACK carries a new cum ack then
   for each TSN  $t_c$  being cum acked for the first time, that was not acked through prior
   gap reports do
     (i) let  $d_c$  be the destination to which  $t_c$  was sent;
     (ii) set  $d_c.find\_pseudo\_cumack = \text{TRUE}$ ;
     (iii) set  $d_c.new\_pseudo\_cumack = \text{TRUE}$ ;
3) if gap reports are present in the SACK then
   for each TSN  $t_p$  being processed from the retransmission queue do
     (i) let  $d_p$  be the destination to which  $t_p$  was sent;
     (ii) if ( $d_p.find\_pseudo\_cumack = \text{TRUE}$ ) and  $t_p$  was not acked in the past then
        $d_p.pseudo\_cumack = t_p$ ;
        $d_p.find\_pseudo\_cumack = \text{FALSE}$ ;
     (iii) if  $t_p$  is acked via gap reports for first time and ( $d_p.pseudo\_cumack = t_p$ ) then
        $d_p.new\_pseudo\_cumack = \text{TRUE}$ ;
        $d_p.find\_pseudo\_cumack = \text{TRUE}$ ;
4) for each destination  $d$  do
   if ( $d.new\_pseudo\_cumack = \text{TRUE}$ ) then update cwnd [1], [2];

```

Fig. 3. Cwnd Update for CMT (CUC) Algorithm – Handling side-effect of reduced cwnd growth due to fewer cwnd updates

we remind the reader that gap reports alone do not (necessarily) imply congestion loss; SACK information is treated only as a concise description of the TSNs received thus far by the receiver.

**Algorithm Details:** Figure 3 shows the proposed Cwnd Update for CMT (CUC) algorithm. A *pseudo-cumack* tracks the earliest outstanding TSN per destination at the sender. An advance in a pseudo-cumack is used by a sender to trigger a cwnd update for the corresponding destination, even when the actual cum ack is not advanced. The pseudo-cumack is used only for cwnd updates; only the actual cum ack can dequeue data in the sender’s retransmission queue since a receiver can renege on data that has been acked through gap reports, but not cumulatively acked. An advantage of CUC is that only the sender’s behavior is affected.

The CUC algorithm introduces three variables per destina-

tion at a sender:

- 1) *pseudo\_cumack* - maintains next pseudo-cumack expected at a sender.
- 2) *new\_pseudo\_cumack* - flag used to indicate if a new pseudo-cumack has been received.
- 3) *find\_pseudo\_cumack* - flag used to find a new pseudo-cumack. This flag is set after a new pseudo-cumack has been received and indicates that the sender should find the next pseudo-cumack expected for the corresponding destination.

In Figure 3, step (1) initializes the variable *new\_pseudo\_cumack*. Step (2) initiates a search for a new *pseudo\_cumack* by setting *find\_pseudo\_cumack* to TRUE for the destinations on which TSNs cumulatively acked were outstanding. Step (2) also triggers a cwnd update by setting

*new\_pseudo\_cumack* to TRUE for those destinations. Step (3) then processes the outstanding TSNs at a sender, and tracks on a per destination basis, the TSN expected to be the next *pseudo\_cumack*. Step (4) finally updates the *cwnd* for a destination if a new pseudo cumack was seen for that destination.

### C. Curbing Increase in Ack Traffic

Sending an ack after receiving every 2 data PDUs (i.e., delayed acks) in SCTP (and TCP) reduces ack traffic in the Internet, thereby saving processing and storage at routers on the ack path. SCTP specifies that a receiver should use the delayed ack algorithm as given in RFC 2581, where acks are delayed only as long as the receiver receives data in order. Reordered PDUs should be acked immediately [5]. With CMT's frequent reordering, this rule causes an SCTP receiver to frequently *not* delay acks. Hence a negative side-effect of reordering with CMT is increased ack traffic.

To prevent this increase, we propose that a CMT receiver ignore the rule mentioned above. That is, a CMT receiver does not immediately ack an out-of-order PDU, but delays the ack. Thus, a CMT receiver always delays acks, irrespective of whether or not data is received in order<sup>2</sup>. Though this modification at the receiver eliminates the increase in ack traffic, RFC 2581's rule has another purpose which gets hampered.

An underlying assumption that pervades SCTP's (and TCP's) design is that data in general arrives in order; data received out-of-order indicates possible loss. According to RFC 2581, "Out-of-order data segments SHOULD be acknowledged immediately, in order to accelerate loss recovery. To trigger the fast retransmit algorithm, the receiver SHOULD send an immediate ... ACK when it receives a data segment above a gap in the sequence space" [5]. In SCTP, four acks with gap reports for a missing TSN (i.e., four missing reports for a TSN) indicate that a receiver received at least four data PDUs sent after the missing TSN. Receipt of four missing reports for a TSN triggers the fast retransmit algorithm at the sender. In other words, the sender has a *reordering threshold* (or *dupack threshold* in TCP terminology) of four PDUs. Since a CMT receiver cannot distinguish between loss and reordering introduced by a CMT sender, the modification suggested above by itself would cause the receiver to delay acks even in the face of loss. When a loss does occur with our modification to a receiver, fast retransmit would be triggered by a CMT sender only after the receiver receives eight(!) data PDUs sent after a lost TSN - an overly conservative behavior.

The effective increase in reordering threshold at a sender can be countered by reducing the actual number of acks required to trigger a fast retransmit at the sender, i.e., by increasing the number of missing reports registered per ack. In other words, if a sender can increment the number of missing reports more accurately per ack received, fewer acks will be required to be

<sup>2</sup>We are concerned with only the case where acks are not delayed due to reordered data. We do not modify a receiver's behavior when an ack being delayed can be piggybacked on reverse path data, or when the delayed ack timer expires.

received by the sender to trigger a fast retransmit. A receiver can provide more information in each ack to assist the sender in accurately inferring the number of missing reports per ack for a lost TSN. We propose that in each ack, a receiver report the count of data PDUs received since the previous ack was sent. A sender then infers the number of missing reports per TSN based on the TSNs being acked in a SACK, number of PDUs reported by the receiver, and knowledge of transmission destination for each TSN. We note that additionally, heuristics (as proposed in [15]) may be used at a CMT sender to address network induced reordering.

**Algorithm Details:** The proposed Delayed Ack for CMT (DAC) algorithm (Figure 4) specifies a receiver's behavior on receipt of data, and also a sender's behavior when the missing report count for a TSN needs to be incremented<sup>3</sup>. Since SCTP (and TCP) acks are cumulative, loss of an ack will result in loss of the data PDU count reported by the receiver, but the TSNs will be acknowledged by the following ack. Receipt of this following ack can cause ambiguity in inferring missing report count per destination. Our algorithm conservatively assumes a single missing report count per destination in such ambiguous cases. The DAC algorithm requires modifications to both the sender and the receiver.

No new variables are introduced in this algorithm, as we build on the SFR algorithm. There is an additional number to be reported in the SACKs for which we use the first 2 bits of the flags field in the SACK chunk header.

In Figure 4, at the receiver side, steps (1) and (2) are self explanatory. The sender side algorithm modifies step (4) of SFR, which determines whether missing report count should be incremented for a TSN. The DAC algorithm dictates *how many* to increment by. Step (4-i) checks if only one destination was newly acked, and allows incrementing missing reports by more than one for TSNs outstanding to that destination. Further, all newly acked TSNs should have been sent later than the missing TSN. If there are newly acked TSNs that were sent before the missing TSN, step (4-i-a) conservatively allows incrementing by only one. If more than one destinations are newly acked, step (4-ii) conservatively allows incrementing by only one.

Figure 5 shows *cwnd* evolution for CMT after including the SFR, CUC and DAC algorithms, i.e.,  $CMT_{scd}$ . With the negative side-effects addressed, we hoped to see  $CMT_{scd}$ 's *cwnd* growth to come close to the expected aggregate *cwnd* growth. In fact, we observed that  $CMT_{scd}$  *cwnd* growth *exceeded* the expected aggregate *cwnd* growth!

To explain this surprising result, we remind the reader that the expected aggregate *cwnd* is the sum of the *cwnd* growth of two independent SCTP runs, each using one of the two destination addresses as its primary destination. In each SCTP run, one delayed ack can increase the *cwnd* by at most one MTU during slow start, even if the ack acks more than one MTU worth of data. On the other hand, we observe in the  $CMT_{scd}$  run that if a delayed ack simultaneously acks an MTU of data on each of the two destinations, the sender

<sup>3</sup>The DAC algorithm can also be used in scenarios where ack traffic lesser than with delayed acks is desirable, such as in data center environments [23].

On receipt of a data PDU [Receiver side behavior]:

- 1) delay sending an ack as given in [2], with the additional change that acks should be delayed even if reordering is observed.
- 2) in each ack, report number of data PDUs received since sending of previous ack.

When incrementing missing report count through SFR:Step (4) (Figure 2) [Sender side behavior]:

- 4) to determine whether missing report count for a TSN  $t_m$  should be incremented:
  - let  $d_m$  be the destination to which  $t_m$  was sent;
    - if** ( $d_m.saw\_newack = \text{TRUE}$ ) **and** ( $d_m.highest\_in\_sack\_for\_dest > t_m$ ) **then**
      - (i) **if** ( $\forall$  destinations  $d_o$  such that  $d_o \neq d_m$ ,  $d_o.saw\_newack = \text{FALSE}$ ) **then**
        - /\*\* all newly acked TSNs were sent to the same destination as  $t_m$  \*\*/
        - (a) **if** ( $\exists$  newly acked TSNs  $t_b, t_a$  such that  $t_b < t_m < t_a$ ) **then**
          - (conservatively) increment missing report count for  $t_m$  by 1;
        - (b) **else if** ( $\forall$  newly acked TSNs  $t_a, t_a > t_m$ ) **then**
          - increment missing report count for  $t_m$  by number of PDUs reported by receiver;
    - (ii) **else**
      - /\*\* Mixed SACK - newly acked TSNs were sent to multiple destinations \*\*/
      - (conservatively) increment missing report count for  $t_m$  by 1;

Fig. 4. Delayed Ack for CMT (DAC) Algorithm – Handling side-effect of increased ack traffic




Fig. 5.  $\text{CMT}_{scd}$ : Evolution of the different cwnds

simultaneously increases each of two cwnds by one MTU. Thus, a single delayed ack in  $\text{CMT}_{scd}$  that acks data flows on two paths causes an aggregate cwnd growth of two MTUs. With delayed acks during slow start, each SCTP association grows its cwnd by 1.5 times each RTT, whereas  $\text{CMT}_{scd}$  increases its cwnd by twice each RTT. We conclude that such delayed acks which simultaneously contribute to the cwnd growth of two destinations helped the aggregate cwnd growth of  $\text{CMT}_{scd}$  to exceed the expected aggregate cwnd growth.

This phenomenon occurs in slow start, therefore benefiting  $\text{CMT}_{scd}$  initially and during some timeout recovery periods. Though the aggregate cwnd growth exceeds expected aggregate cwnd growth, we argue that the sender is not overly aggressive, i.e., not TCP-unfriendly. The sender is able to clock out more data due to delayed acks that acknowledge data flows on multiple paths. The sender does not create bursts of

data during slow start, and builds up the ack clock as expected. Though it does not improve  $\text{CMT}_{scd}$ 's performance significantly, this phenomenon demonstrates a benefit of *sequence space sharing* among flows on different paths that occurs within a  $\text{CMT}_{scd}$  association.

#### IV. CMT PERFORMANCE EVALUATION

With correct behavior ensured by  $\text{CMT}_{scd}$ , we now evaluate its performance<sup>4</sup>. In Section IV-A, we discuss our methodology for evaluating CMT. In Section IV-B, we present five retransmission policies for CMT. In Section IV-C, we identify two modifications that must be made to CMT to accommodate the different retransmission policies. In Section IV-D, we evaluate CMT against AppStripe - our reference application for performance evaluation of CMT. In Section IV-E, we compare and analyze the different retransmission policies to decide upon a recommended policy for CMT.

##### A. Evaluation Methodology

As a reference for performance, we use *AppStripe* - a hypothetical multihome-aware application that achieves the highest throughput achievable by an application that distributes data across multiple SCTP associations (see Figure 6). We emphasize that AppStripe performs idealized scheduling at the application layer, and is not doable in practice<sup>5</sup>. End-to-end load sharing is performed at the application layer by AppStripe, and at the transport layer by CMT [24].

The simulation topology (see Figure 7) is simple - the edge links represent the last hop, and the core links represent end-to-end conditions on the Internet. This simulation topology does not account for effects seen in the Internet and other real networks such as network induced reordering, delay spikes,

<sup>4</sup>Henceforth, we refer to  $\text{CMT}_{scd}$  as simply CMT.

<sup>5</sup>We simulate AppStripe by post-processing simulation traces.

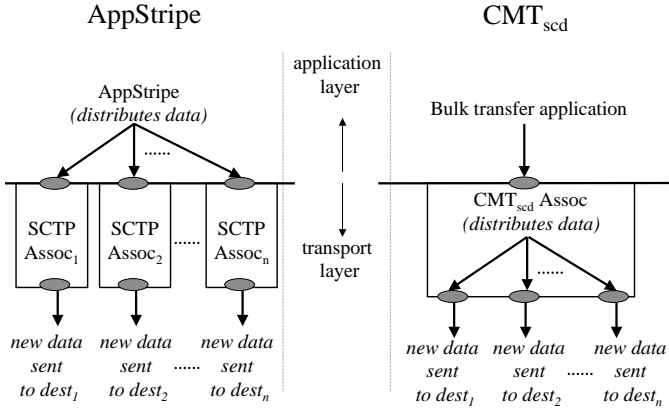


Fig. 6. Schematic - AppStripe and CMT

etc.; these effects are beyond the scope of this study. Our simulation evaluation provides insight into the fundamental differences between AppStripe and CMT, and between the different retransmission policies in a constrained environment. We chose a simple topology to avoid influence of other effects, and to focus on performance differences which we believe should hold true in a real environment as well<sup>6</sup>. The loss rate on Path 1 is maintained at 1%, and on Path 2 is varied from 1 to 10%. A loss rate of 1% means a forward path loss rate of 1%, and a reverse path loss rate of 1%.




Fig. 7. Simulation topology used for evaluation

### B. CMT Retransmission Policies

Multiple paths present an SCTP sender with several choices where to send a retransmission of a lost transmission. But these choices are not well-informed since SCTP restricts sending new data, which can act as probes for information (such as available bandwidth, loss rate and RTT), to only one primary destination. Consequently, an SCTP sender has minimal information about other paths to a receiver. On the other hand, a CMT sender maintains more current and more accurate information about all paths to a receiver, since new

data is regularly being sent to all destinations concurrently. This information allows a CMT sender to make a more informed decision where to send a retransmission. We now investigate how CMT should make this decision in a realistic lossy environment.

We present five retransmission policies for CMT [25]. For four policies, a retransmission may be sent to a destination other than the one used for the original transmission. Previous work on retransmission policies for SCTP [26] shows that sending retransmissions to an alternate destination degrades performance primarily because of the lack of sufficient traffic on alternate paths. With CMT, data is concurrently sent on all paths, thus rendering the results in [26] not applicable. The five different retransmission policies for CMT are:

- **RTX-SAME** - Once a new data chunk is scheduled and sent to a destination, all retransmissions of the chunk thereafter are sent to the same destination (until the destination is deemed *inactive* due to failure [2]).
- **RTX-ASAP** - A retransmission of a data chunk is sent to any destination for which the sender has cwnd space available at the time the retransmission needs to be sent. If the sender has available cwnd space for multiple destinations, one is chosen randomly.
- **RTX-CWND** - A retransmission of a data chunk is sent to the destination for which the sender has the largest cwnd. A tie is broken by random selection.
- **RTX-SSTHRESH** - A retransmission of a data chunk is sent to the destination for which the sender has the largest ssthresh. A tie is broken by random selection.
- **RTX-LOSSRATE** - A retransmission of a data chunk is sent to the destination with the lowest loss rate path. If multiple destinations have the same loss rate, one is selected randomly.

Of the policies, RTX-SAME is simplest. RTX-ASAP is a “hot-potato” retransmission policy - the goal is to retransmit as soon as possible without regard to loss rate. RTX-CWND and RTX-SSTHRESH practically track, and attempt to move retransmissions onto the path with the estimated lowest loss rate. Since ssthresh is a slower moving variable than cwnd, the values of ssthresh may better reflect the conditions of the respective paths. RTX-LOSSRATE uses information about loss rate provided by an “oracle” - information that RTX-CWND and RTX-SSTHRESH estimate. This policy represents a “hypothetically” ideal case; hypothetical since in practice, a sender typically does not know apriori which path has the lowest loss rate; ideal since the path with the lowest loss rate has highest chance of having a packet delivered. We initially hypothesized that retransmission policies that take loss rate into account would outperform ones that do not.

### C. Modifications to Protocol Mechanisms

Two protocol modifications are needed in CMT to allow redirecting retransmissions to a different destination than the original transmission.

1) *CUCv2: Modified CUC Algorithm*: The CUC algorithm (Figure 3) enables correct cwnd updates in the face

<sup>6</sup>The simulation topology is clearly simplistic. Unfortunately, more realistic complex topologies involving variable cross-traffic required too much time to simulate. We believe that the relative performance of the evaluated policies will be the same for our simple topology as well as a more complex one.

```

At beginning of an association [Sender side behavior]:
  ∀ destinations d, reset
    d.find_pseudo_cumack = TRUE;
    d.find_rtx_pseudo_cumack = TRUE;
On receipt of a SACK [Sender side behavior]:
  1) ∀ destinations d, reset
    d.new_pseudo_cumack = FALSE;
    d.new_rtx_pseudo_cumack = FALSE;
  2) if the ack carries a new cum ack then
    for each TSN  $t_c$  being cum acked for the first time, that was not acked through prior
    gap reports do
      (i) let  $d_c$  be the destination to which  $t_c$  was sent;
      (ii) set  $d_c.find\_pseudo\_cumack = TRUE$ ;
      (iii) set  $d_c.new\_pseudo\_cumack = TRUE$ ;
      (iv) set  $d_c.find\_rtx\_pseudo\_cumack = TRUE$ ;
      (v) set  $d_c.new\_rtx\_pseudo\_cumack = TRUE$ ;
  3) if gap reports are present in the ack then
    for each TSN  $t_p$  being processed from the retransmission queue do
      (i) let  $d_p$  be the destination to which  $t_p$  was sent;
      (ii) if ( $d_p.find\_pseudo\_cumack = TRUE$ ) and  $t_p$  was not acked in the past
      and  $t_p$  was not retransmitted then
         $d_p.pseudo\_cumack = t_p$ ;
         $d_p.find\_pseudo\_cumack = FALSE$ ;
      (iii) if  $t_p$  is acked via gap reports for first time and ( $d_p.pseudo\_cumack = t_p$ ) then
         $d_p.new\_pseudo\_cumack = TRUE$ ;
         $d_p.find\_pseudo\_cumack = TRUE$ ;
      (iv) if ( $d_p.find\_rtx\_pseudo\_cumack = TRUE$ ) and  $t_p$  was not acked in the past
      and  $t_p$  was retransmitted then
         $d_p.rtx\_pseudo\_cumack = t_p$ ;
         $d_p.find\_rtx\_pseudo\_cumack = FALSE$ ;
      (v) if  $t_p$  is acked via gap reports for first time and ( $d_p.rtx\_pseudo\_cumack = t_p$ ) then
         $d_p.new\_rtx\_pseudo\_cumack = TRUE$ ;
         $d_p.find\_rtx\_pseudo\_cumack = TRUE$ ;
  4) for each destination  $d$  do
    if ( $d.new\_pseudo\_cumack = TRUE$ ) or ( $d.new\_rtx\_pseudo\_cumack = TRUE$ ) then
      Update cwnd [1], [2];

```

Fig. 8. CUCv2 Algorithm - Modified Cwnd Update for CMT (CUC) Algorithm

of increased reordering due to CMT. To recap, this algorithm recognizes a set of TSNs outstanding per destination, and the per-destination *pseudo\_cumack* traces the left edge of this list of TSNs, per destination. The CUC algorithm though, assumes that retransmissions are sent to the same destination as the original transmission. The per-destination *pseudo\_cumack* therefore moves whenever the corresponding left edge is acked; the TSN on the left edge being acked may or may not have been retransmitted.

If the assumption about the retransmission destination is violated, and a retransmission is made to a different destination from the original, the CUC algorithm cannot faithfully track the left edge on either destination. We modify the CUC algorithm to permit the different retransmission policies. The modified CUC algorithm, named CUCv2 is shown in Figure 8.

The crux of the modification is in recognizing that of the TSNs outstanding on a destination, a distinction can be made -

those that have been retransmitted, and those that have not. The CUCv2 algorithm maintains two left edges for these two sets of TSNs - *rtx-pseudo-cumack* and *pseudo-cumack*. Whenever either of the left edges moves, a cwnd update is triggered. In CUCv2, lines 2(iv), 2(v), 3(iv) and 3(v) have been added, and lines 3(ii) and 4 have been modified from the CUC algorithm.

2) *Spurious Timeout Retransmissions*: When a timeout occurs, an SCTP sender is expected to bundle and send as many of the earliest TSNs outstanding on the destination for which the timeout occurred as can fit in an MSS (Maximum Segment Size) PDU. As per RFC 2960, if more TSNs are outstanding on that destination, these TSNs “should be marked for retransmission and sent as soon as cwnd allows (normally when a SACK arrives)”. This rule is intuitive - while sending, retransmissions are generally given priority over new transmissions. As in TCP, the cwnd is also collapsed to 1 MSS for the destination on which a timeout occurs, allowing only one



MSS sized PDU in flight.

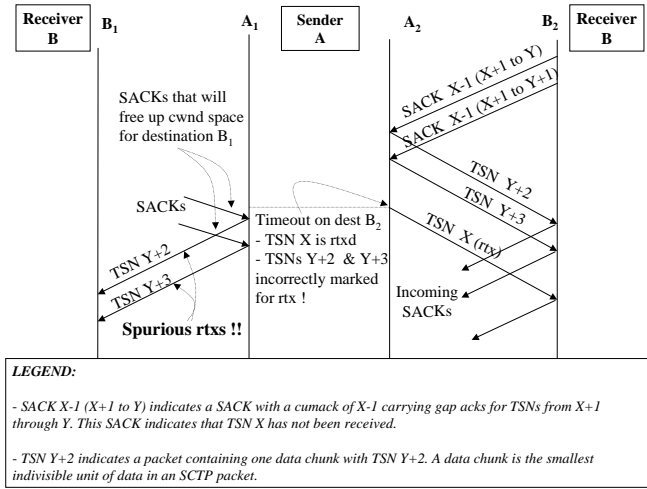


Fig. 9. Example of spurious retransmissions after timeout in CMT

A timeout retransmission can occur in SCTP (as in TCP) for several reasons. One reason is loss of the fast retransmission of a TSN<sup>7</sup>. Consider Figure 9. When a timeout occurs due to loss of a fast retransmission, some TSNs that were just sent to the destination on which the timeout occurred are likely awaiting acks (in Figure 9, TSNs  $Y+2$  and  $Y+3$ )<sup>8</sup>. These TSNs get incorrectly marked for retransmission on timeout. With the different retransmission policies in CMT, the retransmissions may be sent to a different destination than the original transmission; if cwnd space for a destination is available, possibly due to receipt of an ack from that destination, TSNs marked for retransmission may be sent to that destination. In Figure 9, spurious retransmissions of TSNs  $Y+2$  and  $Y+3$  are sent to destination  $B_1$ , on receipt of acks freeing up cwnd space for destination  $B_1$ . Spurious retransmissions are exacerbated in CMT, as shown through this illustration, due to the possibility of sending data (including retransmissions) to multiple destinations concurrently.

We simulated the occurrence of such spurious retransmissions with the different retransmission policies in CMT. The simulation topology used is the one described in Section IV-A. Figure 10(a) shows the ratio of retransmissions relative to the number of actual packet drops at the router. Ideally, the two numbers should be equal; all curves should be straight lines at  $y = 1$ . Figure 10(a) shows that spurious retransmissions occur commonly in CMT with the different retransmission policies.

We propose a heuristic to avoid these spurious retransmissions. Our heuristic assumes that a timeout cannot be triggered on a TSN until the TSN has been outstanding for at least one Round-Trip Time (RTT). Thus, if a timeout is triggered, TSNs which were sent within an RTT in the past are not marked for

retransmission. We use an average measure of the RTT for this purpose - the smoothed RTT, which is maintained at a sender. This heuristic requires the sender to maintain a timestamp for each TSN indicating the time at which the TSN was last transmitted (or retransmitted). Figure 10(b) shows how the application of this heuristic drastically reduces the number of spurious retransmissions for all retransmission policies.

#### D. Performance of CMT vs. AppStripe

Figure 11 compares the time taken to transfer an 8MB file using CMT with the five retransmission policies, vs. using AppStripe. The x-axis represents different loss rates on Path 2. Each plotted value is the mean of at least 30 simulation runs. Overall, AppStripe ( $\times$  in Figure 11) performs worst, and CMT using any of the retransmission policies performs better than AppStripe; some policies better than others. At a 7% loss rate on Path 2, AppStripe takes 40.4 seconds to transfer an 8 MB file, whereas CMT using RTX-SAME or RTX-CWND takes 35.5 or 33.2 seconds, respectively. We first discuss the performance difference between CMT in general and AppStripe.

CMT using any retransmission policy performs better than AppStripe, particularly as the loss rate on Path 2 increases. Note that our AppStripe represents the *best possible performance* expected by an application that stripes data over multiple SCTP associations. AppStripe is an idealized case; CMT's performance gain over a practical AppStripe implementation would be even larger. A practical implementation has to also address issues such as striping data across paths that have different and changing delays and loss rates. Such issues may require information from the transport layer (such as current cwnd and RTT), that may not be readily available to the application.

CMT performs better than AppStripe for two reasons. First, and more significant, CMT is more resilient to reverse path loss than AppStripe. CMT uses a single sequence space (TSN space, used for congestion control and loss detection and recovery) across an association's multiple paths, whereas AppStripe by design uses an independent sequence space per path. Since acks are cumulative, sharing of sequence spaces across paths helps a CMT sender receive ack info on either of the return paths. Thus, CMT effectively uses *both* return paths for communicating ack info to the sender, whereas each association in AppStripe cannot help the other "ack-wise". These results demonstrate the significant result that CMT's sharing of sequence space across paths is an *inherent benefit that performing load sharing at the transport layer has over performing it at the application layer*.

Second, CMT gets faster overall cwnd growth than AppStripe in slow start (See Section III-C). As loss increases, number of timeouts increases, and since the sender enters slow start after each timeout, the sender spends more time overall in slow start.

#### E. Performance of different retransmission policies for CMT

Of the retransmission policies for CMT in Figure 11, RTX-SAME ( $\ominus$ ) performs marginally but consistently worse than

<sup>7</sup>The Multiple Fast Retransmit (MFR) algorithm allows recovery using fast retransmission multiple times on the same TSN [27], but has not been ported to CMT. The only recovery mechanism from the loss of a fast retransmission in CMT, as is the case currently in SCTP and TCP, is a timeout recovery.

<sup>8</sup>This figure illustrates the point. Discrepancies in scale may be overlooked.




Fig. 10. Spurious retransmissions in CMT: (a) Without RTT heuristic (b) With RTT heuristic

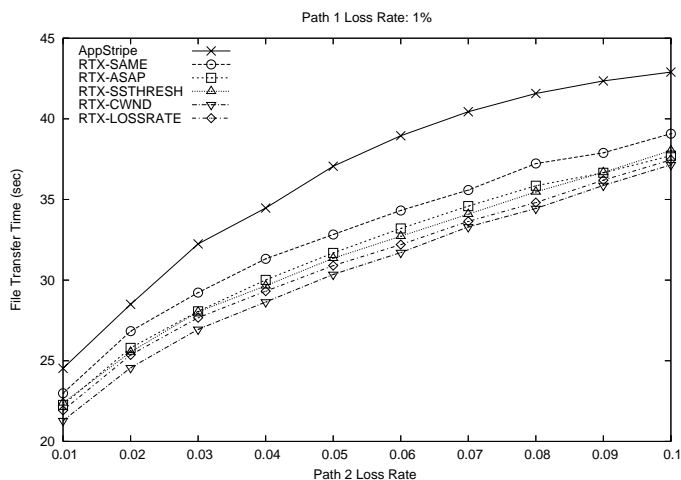


Fig. 11. With Path 1 loss rate = 1%, performance of AppStripe vs. CMT with different retransmission policies

RTX-ASAP ( $\square$ ), which in turn performs almost as well as the loss rate based policies - RTX-SSTHRESH ( $\triangle$ ), RTX-CWND ( $\nabla$ ), and RTX-LOSSRATE ( $\diamond$ ). Of the loss rate based policies, RTX-CWND outperforms RTX-SSTHRESH and RTX-LOSSRATE, although marginally. While the performance difference between the retransmission policies in Figure 11 are not significant, these results use an 8MB receiver's buffer (rbuf) that does not constrain the sender - an unrealistic assumption which we will now drop [28].

Figure 12(a) shows the time taken for a CMT sender to transfer an 8MB file when the rbuf is set to 64KB, using the five retransmission policies. RTX-SAME is the simplest to implement, but performs worst. The performance difference between RTX-SAME and other policies increases as the loss rate on Path 2 increases. RTX-ASAP performs better than RTX-SAME, but still worse than RTX-LOSSRATE, RTX-SSTHRESH and RTX-CWND. The three loss rate based policies perform equally.

Figure 12(b) shows the number of retransmission timeouts experienced when using the different policies. This figure shows that performance improvement in using RTX-LOSSRATE, RTX-CWND, and RTX-SSTHRESH is due to the reduced number of timeouts. A lost transmission may be recovered via a fast retransmission, but a lost fast retransmission can be recovered only through a timeout. RTX-SAME does not consider loss rate in choosing a retransmission destination and consequently experiences the largest number of timeouts due to increased loss of retransmissions.

RTX-ASAP does not consider loss rate, and performs better than RTX-SAME. This improved performance with RTX-ASAP is attributed to cwnd space availability on both destinations most of the times a retransmission is triggered - (i) one retransmission is normally allowed to be sent to the destination that has just suffered loss, and (ii) the ack that triggers a retransmission (in case of fast retransmission) may have created cwnd space for the other destination. From (i) and (ii), RTX-ASAP has cwnd space availability on both destinations to send a retransmission. Consequently, RTX-ASAP randomly chooses a destination causing a reduction in timeouts over RTX-SAME which pins its TSNs to the same destination. The three loss rate based policies effectively choose the better destination to redirect retransmissions to, and thus show fewer timeouts than RTX-ASAP.

Figures 13(a) and (b) show performance of the retransmission policies with rbuf sizes of 32KB and 128KB respectively. Together with Figure 12(a), we can see that the smaller the rbuf, the more important the choice of retransmission policy. These results show that *a retransmission policy that considers loss outperforms policies that do not, particularly in the practical reality where rbuf is constrained.*

Figures 12 and 13 suggest that any retransmission policy that takes loss rate into account will likely improve load distribution for both new transmissions and retransmissions. Retransmissions will be redirected to a lower loss rate path, avoiding inactive timeout recovery periods, and allowing new transmissions to be sent on the higher loss rate path, thus




Fig. 12. With  $rbuf = 64KB$ , and Path 1 loss rate = 1%.: (a) Time taken by CMT to transfer an 8MB file, (b) Number of retransmission timeouts for CMT with different retransmission policies




Fig. 13. With Path 1 loss rate = 1%, time taken to transfer an 8MB file using: (a)  $rbuf=32KB$ , (b)  $rbuf=128KB$

maintaining a flow of data on both paths. Policies that take loss rate into account avoid repeated retransmissions and timeouts - thus also improving the timeliness of data.

Of three loss rate based policies, the practical ones to implement are RTX-CWND and RTX-SSTHRESH. Both perform equally under all conditions considered. *sssthresh*, being a slower moving variable, may better reflect the condition of a path. *We therefore select RTX-SSTHRESH as the recommended retransmission policy for CMT.*

## V. SUMMARY AND DISCUSSION

We identified three negative side-effects of introducing CMT with SCTP, and proposed algorithms to avoid these side-effects. We compared CMT against AppStripe, an idealized data striping application, and showed that *a shared sequence space in CMT improves performance and increases resilience to reverse path loss*. We also presented and evaluated five retransmission policies for CMT. Our results reveal that a

retransmission policy that considers loss rate performs better than one that does not, particularly in the practical reality where  $rbuf$  is constrained. *We recommend the RTX-SSTHRESH retransmission policy for CMT.*

CMT also inherently adds to SCTP's fault tolerance, which is a major motivation for, and benefit of multihoming. An SCTP sender sends data (which act as implicit probes) to a primary destination, and gathers information about paths to all other destination addresses (or *alternate paths*) through explicit probes. Since explicit probes are infrequent, a sender usually has inadequate information about alternate paths to a receiver. Consequently, a sender is unable to make informed decisions about a new destination to use when the primary destination is unreachable due to a network failure. A CMT sender avoids this problem because data sent concurrently on all paths act as frequent implicit probes, reflecting current conditions of paths to all destination addresses. This information will better assist a CMT sender in detecting and responding

to network failure events.

We emphasize that though CMT uses SCTP, our goal was to study CMT at the transport layer in general. The reordering issues and solutions (Section III), and the retransmission policies and their evaluation (Section IV) should be relevant even if CMT used a different transport protocol.

## VI. RELATED WORK

This section discusses load balancing efforts at the application (Section VI-A), transport (Section VI-B) and network (Section VI-C) layers, a classification also suggested in [29].

### A. Load Balancing at the Application Layer

Several applications have been proposed to use multiple concurrent transport layer connections to a receiver. Some proposals [30], [31] use multiple TCP connections to increase throughput in high bandwidth networks. These applications load balance over the same path to a receiver, whereas CMT distributes data over multiple independent paths to a receiver.

Mao et al. [32] extend RTP (Realtime Transport Protocol) to support use of multiple paths in *Multi-path Realtime Transport Protocol (MRTP)*, an application layer protocol which could use any of TCP, SCTP or UDP as transport. MRTP specifies session establishment and maintenance mechanisms and scheduling mechanisms over multiple paths. The authors propose, as one option, to use SCTP multihoming for simultaneously using multiple paths. [32] is complementary to CMT work, since it provides motivation and an application that would benefit from using CMT in a multipath environment.

Content Networks [33], [34] provide an infrastructure for *connection level load balancing* at the granularity of TCP connections. Connection level load balancing is useful for short TCP connections such as web requests and responses, but can be suboptimal for long transfers such as bulk data transfers, where the server is constrained to a single path throughout the transfer. In CMT, we consider load balancing *within* a transport connection.

Load balancing at the transport layer is desirable since the transport layer, being the lowest end-to-end layer, has the most accurate information about the end-to-end path(s). CMT maintains loss and delay information which are of significant value when redirection of retransmissions needs to be made - such decisions are best made in the transport layer. Besides the performance benefits with CMT discussed in Section IV-D, we believe that load balancing at the application layer increases code redundancy and room for error by requiring independent implementations in each application rather than coded once in the transport layer.

### B. Load Balancing at the Transport Layer

Hsieh et al. [29] propose a transport protocol called *pTCP (parallel TCP)* which provides an infrastructure for data stripping within the transport layer. pTCP has two components - Striped connection Manager (SM) and TCP-virtual (TCP-v). The TCP-v's are separate connections that are managed by the SM. TCP-v probes the path and performs congestion

control and loss detection/recovery, while the SM decides which data is sent on which TCP-v. This important decoupling of functionality avoids some pitfalls of the application layer approaches, and allows for intelligent scheduling of transmissions and retransmissions. A significant issue with pTCP is its complexity. As the authors note themselves, maintenance of multiple Transmission Control Blocks (TCBs) at a sender can be a resource sink [35]. The complexity of implementation is also non-trivial, since pTCP replicates a lot of transport layer functionality such as connection establishment/teardown and checksum calculations. pTCP has several unresolved issues as well. For instance, if both sender and receiver are multihomed with two IP addresses each, pTCP does not address how a sender decides which sender-receiver pairs to establish TCP connections on - a complex problem. Plugging transport protocols into pTCP also requires non-trivial modifications to the transport protocols themselves. CMT, on the other hand, uses and proposes modifications to SCTP, a transport protocol which already has built-in mechanisms for handling multihoming.

Al et al. [36] suggest ideas for *load sharing with SCTP*. Their solutions require additional metadata in the SCTP PDUs. We believe that the SCTP (and TCP-SACK) PDUs already contain sufficient information for the data sender to infer the per-path ordering information that [36] explicitly codes as metadata. [36] fails to suggest modified procedures for mechanisms which are immediately affected, such as initialization of the per-path sequence numbers, association initialization and shutdown procedures with multiple sequence numbering schemes, and response to renegeing by a receiver. We have also seen that sharing sequence number space across paths improves performance - [36] uses a separate sequence number space per path, and will therefore not see CMT's performance benefits. Further, [36] assumes that the rbuf does not constrain a sender which is unrealistic in practice.

Argyriou et al. [37] provide techniques for *bandwidth aggregation with SCTP*, but do not present and analyze their protocol modifications to SCTP. The modification to the fast retransmission algorithm that is presented is simplistic and assumes information that will likely not be available to an SCTP receiver. For instance, the implicit assumption that a receiver will be able to differentiate a packet loss from reordering is a strong and unrealistic assumption. [37] also ignores the impact of a bounded rbuf.

### C. Load Balancing at the Network Layer

Phatak and Goff [38] propose distributing data at the network (IP) layer transparent to the higher layers using IP-in-IP encapsulation. The authors identify conditions under which this mechanism would work without triggering incorrect retransmission timeouts. [38] assumes that end-to-end delays are dominated by fixed transmission delay, and fails to address propagation delay dominated paths, and paths with dynamically changing bandwidths and delays.

Several proposals exist for *multipath routing* - routing packets from a source to a destination network over multiple paths [39]–[41]. However, different paths are likely to exhibit

different RTTs, thus introducing packet reordering. TCP's performance degrades in the presence of increased reordering. To enable optimal load balancing at intermediate routers without affecting end-to-end TCP performance, modifications to TCP have also been proposed [15]–[17], [42]. These proposals augment and/or modify TCP's congestion control mechanisms to cope with reordering introduced by network layer load balancing; the burden of actually using multiple paths in the network is left to the intermediate routers.

In the Internet, the end user has knowledge of, and control over, only the multihomed end hosts, not the intermediate routers. In such cases the end host cannot dictate or govern use of multiple paths in the network. But the end host can use multiple end-to-end paths available to the host [43], [44], thus motivating CMT at the transport layer.

### DISCLAIMER

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.

### REFERENCES

- [1] R. Stewart, I. Arias-Rodriguez, K. Poon, A. Caro, and M. Tuexen, "Stream Control Transmission Protocol (SCTP) Implementer's Guide," draft-ietf-tsvwg-sctpimpguide-12.txt, Oct. 2004, (work in progress).
- [2] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and V. Paxson, "Stream Control Transmission Protocol," RFC2960, Oct. 2000.
- [3] E. Kohler, M. Handley, and S. Floyd, "Designing DCCP: Congestion Control Without Reliability," Tech. Rep., ICIR, 2004.
- [4] S. Floyd, J. Mahdavi, M. Mathis, and M. Podolsky, "An Extension to the Selective Acknowledgement (SACK) Option for TCP," RFC2883, IETF, July 2000.
- [5] M. Allman, V. Paxson, and W. Stevens, "TCP Congestion Control," RFC2581, IETF, Apr. 1999.
- [6] E. Blanton, M. Allman, K. Fall, and L. Wang, "A Conservative Selective Acknowledgment (SACK) Based Loss Recovery Algorithm For TCP," RFC3517, IETF, Apr. 2003.
- [7] J. Iyengar, *End-to-end Load Balancing using Transport Layer Multihoming*, PhD Dissertation, CISC Dept, University of Delaware, (in progress).
- [8] M. S. Kim, T. Kim, Y. Shin, S. S. Lam, and E. J. Powers, "A Wavelet-based Approach to Detect Shared Congestion," in *ACM SIGCOMM*, Portland, Oregon, Aug. 2004.
- [9] D. Rubenstein, J. Kurose, and D. Towsley, "Detecting Shared Congestion of Flows Via End-to-End Measurement," *IEEE/ACM Transactions on Networking*, vol. 10, no. 3, June 2002.
- [10] D. Katabi, I. Bazzi, and X. Yang, "A Passive Approach for Detecting Shared Bottlenecks," in *IEEE ICCN 2001*, Oct. 2001.
- [11] K. Harfoush, A. Bestavros, and J. Byers, "Robust Identification of Shared Losses Using End-to-End Unicast Probes," in *ICNP 2000*, Osaka, Japan, Oct. 2000.
- [12] UC Berkeley, LBL, USC/ISI, and Xerox Parc, "ns-2 documentation and software," Version 2.1b8, 2001, <http://www.isi.edu/nsnam/ns>.
- [13] A. Caro and J. Iyengar, "ns-2 SCTP module," Version 3.2, December 2002, <http://pe1.cis.udel.edu>.
- [14] J. C. R. Bennett, C. Partridge, and N. Shectman, "Packet Reordering is Not Pathological Network Behavior," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, Dec. 1999.
- [15] E. Blanton and M. Allman, "On Making TCP More Robust to Packet Reordering," *ACM Computer Communication Review*, vol. 32, no. 1, Jan. 2002.
- [16] M. Zhang, B. Karp, S. Floyd, and L. Peterson, "RR-TCP: A Reordering-Robust TCP with DSACK," Tech. Rep. TR-02-006, ICSI, July 2002.
- [17] S. Bohacek, J. P. Hespanha, J. Lee, C. Lim, and K. Obraczka, "TCP-PR: TCP for Persistent Packet Reordering," in *IEEE ICDCS 2003*, Rhode Island, May 2003.
- [18] J. Iyengar, K. Shah, P. Amer, and R. Stewart, "Concurrent Multipath Transfer Using SCTP Multihoming," in *SPECTS 2004*, San Jose, California, July 2004.
- [19] J. Iyengar, A. Caro, P. Amer, G. Heinz, and R. Stewart, "Making SCTP More Robust to Changeover," in *SPECTS 2003*, Montreal, Canada, July 2003.
- [20] J. Iyengar, A. Caro, P. Amer, G. Heinz, and R. Stewart, "SCTP Congestion Window Overgrowth During Changeover," in *SCI 2002*, Orlando, FL, July 2002.
- [21] R. Ludwig and R. Katz, "The Eifel Algorithm: Making TCP Robust Against Spurious Retransmissions," in *ACM Computer Communications Review*, Jan. 2000.
- [22] S. Ladha, S. Baucke, R. Ludwig, and P. Amer, "On Making SCTP Robust to Spurious Retransmissions," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, pp. 123–135, Apr. 2004.
- [23] N. Jani and Krishna Kant, "SCTP Performance in Data Center Environments," Tech. Rep., Intel Corporation, 2005.
- [24] J. Iyengar, P. Amer, and R. Stewart, "Concurrent Multipath Transfer Using Transport Layer Multihoming: Performance Under Varying Bandwidth Proportions," in *MILCOM 2004*, Monterey, CA, Oct. 2004.
- [25] J. Iyengar, P. Amer, and R. Stewart, "Retransmission Policies For Concurrent Multipath Transfer Using SCTP Multihoming," in *ICON 2004*, Singapore, Nov. 2004.
- [26] A. Caro, P. Amer, and R. Stewart, "Transport Layer Multihoming for Fault Tolerance in FCS Networks," in *MILCOM 2003*, Boston, MA, Oct. 2003.
- [27] A. Caro, P. Amer, J. Iyengar, and R. Stewart, "Retransmission Policies with Transport Layer Multihoming," in *ICON 2003*, Sydney, Australia, Sept. 2003.
- [28] J. Iyengar, P. Amer, and R. Stewart, "Receive Buffer Management For Concurrent Multipath Transport Using SCTP Multihoming," Tech Report TR2005-10, CIS Dept, University of Delaware, Jan. 2005.
- [29] H.Y. Hsieh and R. Sivakumar, "A Transport Layer Approach for Achieving Aggregate Bandwidths on Multihomed Mobile Hosts," *ACM/Kluwer Mobile Networks and Applications Journal (MONET)*, vol. 11, no. 1, Jan. 2005.
- [30] T. Hacker and B. Athey, "The End-to-End Performance Effects of Parallel TCP Sockets on a Lossy Wide-Area Network," in *IEEE IPDPS*, Ft. Lauderdale, FL, Apr. 2002.
- [31] H. Sivakumar, S. Bailey, and R. Grossman, "PSockets: The Case For Application-Level Network Striping For Data Intensive Applications Using High Speed Wide Area Networks," in *IEEE Supercomputing (SC)*, Dallas, TX, Nov. 2000.
- [32] S. Mao, D. Bushmitch, S. Narayanan, and S. S. Panwar, "MRTP: A Multi-Flow Realtime Transport Protocol for Ad Hoc Networks," in *IEEE Vehicular Technology Conference*, Orlando, Florida, Oct. 2003.
- [33] M. Day, B. Cain, G. Tomlinson, and P. Rzewski, "A Model For Content Internetworking (CDI)," RFC3466, IETF, Feb. 2003.
- [34] A. Barbir, B. Cain, R. Nair, and O. Spatscheck, "Known Content Network (CN) Request-Routing Mechanisms," RFC3568, IETF, July 2003.
- [35] H.Y. Hsieh and R. Sivakumar, "A Transport Layer Approach for Achieving Aggregate Bandwidths on Multihomed Mobile Hosts," in *ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Atlanta, Georgia, Sept. 2002.
- [36] A. Abd El Al, T. Saadawi, and M. Lee, "LS-SCTP: A Bandwidth Aggregation Technique For Stream Control Transmission Protocol," *Computer Communications*, vol. 27, no. 10, 2004.
- [37] A. Argyriou and V. Madiseti, "Bandwidth Aggregation With SCTP," in *IEEE Globecom 2003*, San Francisco, CA, Dec. 2003.
- [38] D. S. Phatak and T. Goff, "A Novel Mechanism for Data Streaming Across Multiple IP Links for Improving Throughput and Reliability in Mobile Environments," in *IEEE INFOCOM 2002*, New York, NY, June 2002.
- [39] R. Krishnan and J. A. Silvester, "Choice of Allocation Granularity in Multi-Path Source Routing Schemes," in *IEEE INFOCOM*, 1993.
- [40] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, "Highly-Resilient, Energy-Efficient Multipath Routing in Wireless Sensor Networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 4, pp. 11–25, Oct. 2001.
- [41] H. Suzuki and F. A. Tobagi, "Fast Bandwidth Reservation Scheme With Multilink And Multipath Routing In ATM Networks," in *IEEE INFOCOM*, 1992.
- [42] M. Gerla, S. S. Lee, and G. Pau, "TCP Westwood Simulation Studies in Multiple-Path Cases," in *SPECTS 2002*, San Diego, California, July 2002.

- [43] R. Teixeira, K. Marzullo, S. Savage, and G.M. Voelker, "Characterizing and Measuring Path Diversity in Internet Topologies," in *ACM SIGMETRICS (extended abstract)*, San Diego, CA, June 2003.
- [44] R. Teixeira, K. Marzullo, S. Savage, and G.M. Voelker, "In Search of Path Diversity in ISP Networks," in *USENIX/ACM Internet Measurement Conference*, Miami, FL, Oct. 2003.