Counting preimages of TCP reordering patterns

Anders Hansson [a], Gabriel Istrate [b,*]

[a] *Information Sciences (CCS-3), Los Alamos National Laboratory, P.O. Box 1663, Mail Stop M997, Los Alamos, NM 87545, USA*

[b] *eAustria Research Institute, Bd. V. Pârvan no. 4, cam. 045B, Timişoara, RO-300223, Romania*

ARTICLE INFO

ABSTRACT

Packet reordering is an important property of network traffic that should be captured by analytical models of the Transmission Control Protocol (TCP). We study a combinatorial problem motivated by Restored [G. Istrate, A. Hansson, S. Thulasidasan, M. Marathe, C. Barrett, Semantic compression of TCP traces, in: F. Boavida (Ed.), Proceedings of the Fifth IFIP NETWORKING Conference, in: Lecture Notes in Computer Science, vol. 3976, Springer-Verlag, 2006, pp. 123–135], a TCP modeling methodology that incorporates information about packet dynamics. A significant component of this model is a many-to-one mapping B that transforms sequences of packet IDs into *buffer sequences* in a manner that is compatible with TCP semantics. We obtain the following results:

- We give an easy necessary and sufficient condition for an input sequence W to be *valid* (i.e. $A \in B^{-1}(W)$ for some permutation A of $\{1, 2, \ldots, n\}$), and a linear time algorithm that, given a valid buffer sequence W of length n, constructs a permutation A in the preimage of W.
- We show that the problem of counting the number of permutations in $B^{-1}(W)$ has a polynomial time algorithm.
- We also show how to extend these results to sequences of IDs that contain repeated packets.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Consider a sequence of TCP *packets*, identified by their integer IDs, as handled by their *receiver*. The receiver must forward the packet sequence to an *application*, subject to respecting *packet sequence integrity*. That is, at every moment the IDs of packets forwarded to the application must form a contiguous sequence $1, 2, \ldots, m$, for some $m \geq 1$. Packets can arrive out of order and thus need to be buffered. Several copies of a packet can arrive, but only one copy of a given packet is useful (and will be stored, if needed). We assume that the receiver evicts a given packet from the buffer and passes it to the application as soon as possible, i.e., as soon as the packet sequence integrity constraint is satisfied.

A given sequence $A = (A_1, \ldots, A_n)$ of packet IDs yields a corresponding sequence $B(A) = (B_{A,1}, \ldots, B_{A,n})$ representing the evolution of the buffer size. In this paper we are interested in the following problem: given a sequence of positive integers W, what is the complexity of:

(1) Deciding whether there exists a permutation A with $W = B(A)$?
(2) Counting the number of permutations in the set $B^{-1}(W)$?

We will give linear (respectively polynomial) time algorithms for these problems.

---

* Corresponding author. Fax: +40 256 244834.
E-mail addresses: hansson@lanl.gov (A. Hansson), gabrielistrate@acm.org (G. Istrate).

2. Motivation

The problem we described in the introduction arises in the context of analytical modeling of TCP dynamics. Therefore, the reader only interested in the combinatorial aspects of the problem can focus on the remaining sections. This section explains in detail the motivation for the problem.

While a lot of attention has been given to modeling the temporal aspects of TCP traffic (see e.g. Jaiswal et al. [2]), the dynamics of packet IDs has not received the same attention. As Bennett et al. [3] have shown, packet reordering is more widespread than originally believed, and is increasingly becoming so, due to technological advances such as link striping and mobile communications. Packet reordering has many severe effects on overall traffic characteristics; hence it is an important component of TCP dynamics (we refer the reader to [3] for further discussion).

Paper [1] introduced RESTORED, a methodology for semantic compression and regeneration of large TCP traces. RESTORED is based on the following observation: TCP guarantees to deliver an ordered packet stream to the application layer and needs to buffer packets that arrive out of order. Consequently, the received packets can be classified into two types: those that could be immediately passed to the application layer, and those that have to be temporarily buffered. A received packet that allows the buffer to flush is called a *pivot packet*. All packets appearing in order are trivially pivots. RESTORED divides the received sequence into segments, bounded by pivot packets. Segments correspond to one of two *phases*:

- An *ordered phase*, in which no reordering is present; thus there is no need for buffering.
- An *unordered phase*, in which there is reordering and buffering.[1] Each occurrence of this phase ends when a pivot packet is received.

RESTORED preserves packet reordering properties of TCP traffic, up to a notion of semantic equivalence of packet traces. This notion is called *behavioral equivalence* and can be motivated as follows:

Definition 1. Let ACK_i be defined as the smallest integer that does not appear among the first i packet IDs (also, define $ACK_0 = 1$). Parameter ACK_i is called *the acknowledgement* (ACK) *at stage i*.

The previous definition relies on the simplifying assumption that in the implementation of TCP each received packet is ACKed, and that value ACK_i is the only information carried by the ACK packet. Of course, real-life acknowledgment policies of TCP can be more complicated [4].

Consider now the following two packet ID sequences: 4 2 3 1 and 4 3 2 1.

These two sequences trigger identical ACK responses, namely 1 1 1 5, i.e., we arrive at the following two mappings:

$$
\begin{aligned}
4 \quad 2 \quad 3 \quad 1 \quad &\rightarrow \quad 1 \quad 1 \quad 1 \quad 5, \\
4 \quad 3 \quad 2 \quad 1 \quad &\rightarrow \quad 1 \quad 1 \quad 1 \quad 5.
\end{aligned}
\tag{1}
$$

Since TCP is a *receiver-driven* protocol, assuming identical network conditions, and discounting possible differences in the value of the congestion window at the beginning of the sequences, *the two ID sequences trigger identical responses from the receiver, and should thus be regarded as indistinguishable from the standpoint of* TCP *dynamics*.

Definition 2. Two sequences of packets P and Q are *behaviorally equivalent* (written $P \equiv_{beh} Q$) if they lead to the same sequences of ACKs.

In practice one might want a notion of equivalence that is even more restrictive than behavioral equivalence. This was, for instance, the case of RESTORED. Its original motivation was to provide a way to compress TCP traces and estimate various measures of quality of service of the original traces by reconstructing "compatible" sequences. Many measures of packet reordering have been proposed in the networking literature [5–7]. Given such a measure M, one way to guarantee that sequences produced by RESTORED resemble the original sequence with respect to measure M is:

(1) Identify an equivalence notion of ID sequences $\equiv$ such that *M is consistent with respect to* $\equiv$, that is

$$
(\forall A, B): (A \equiv B) \Rightarrow (M(A) = M(B)).
\tag{2}
$$

(2) Make sure that for any sequence A, the sequence $R(A)$ regenerated by RESTORED satisfies $R(A) \equiv A$.

(See also [8] for more discussion and clarification.) Behavioral equivalence might be too coarse (as an equivalence relation) to guarantee consistency of many reordering metrics and, thus, needs to be refined. In a companion paper [9] we have considered such an equivalence notion, based on the following notion of buffer size:

---

[1] A technical assumption we will employ is that duplicates of packets that have already been uploaded to the application layer are discarded. This is a sensible assumption, given TCP behavior.

Definition 3. Let $A = \{A_1, A_2, \ldots, A_n\}$ be a sequence of packet IDs. We define the FB as an operator that after receiving a packet A_i at time index i, outputs the difference between the highest ID (H_i) seen so far and the highest ID (L_i) that could be uploaded.

$$FB(A_i) = H_i - L_i. \tag{3}$$

In other words, FB is the size of the smallest buffer large enough to store all packets that arrive out of order, where the definition of size accounts for reserving space for unreceived packets with intermediate IDs as well. The *buffer sequence* FB(P) associated with a sequence P of packet IDs is simply a time series of FB values computed after each packet has been received.

Two sequences of packet IDs P and Q are *FB equivalent* (written $P \equiv_{FB} Q$) if FB(P) = FB(Q).

This definition is directly related to the semantics of TCP, since it preserves quantities such as the size of the receiving window, as measured by parameter *RcvWindow* (see [10], pp. 246–247). Inverting the mapping FB can be done in polynomial time [9]. However, the complexity of computing the cardinality of the preimage $FB^{-1}(W)$ was left open, and was only solved in two special cases.

In this paper, we use a different notion, introduced below, for which more precise results can be obtained.

Definition 4. Buffer size is the smallest size of a buffer that can store all out-of-order packets. Two sequences of packets P and Q are *buffer equivalent* (written $P \equiv_{buf} Q$) if $B(P) = B(Q)$, that is the sequences of buffer sizes associated with receiving P and Q are identical.

From a combinatorial perspective, buffer equivalence is more natural than FB equivalence. Its relation with behavioral equivalence is, however, slightly more complicated:

(1) Buffer equivalence is *not* a refinement of behavioral equivalence in general. Indeed, sequences of packet IDs 2 3 3 1 and 3 4 1 2 are buffer equivalent (they both map to sequence 1 2 2 0) but *not* behaviorally equivalent (the ACKs are 1 1 1 4 and 1 1 2 5, respectively). This stands in contrast to FB equivalence which is indeed [9] a refinement of behavioral equivalence.
(2) Buffer equivalence refines behavioral equivalence when restricted to *permutations* (sequences with no repeats or lost packets). For a formal statement and proof of this claim see Proposition 1.
(3) Finally, buffer equivalence is incomparable (as an equivalence notion) with FB equivalence [8].

On the other hand there exist reordering metrics M defined in the networking literature (e.g. *reorder buffer density* [11]) with the following properties:

(1) M only depends on packets received for the first time, and not on repeat packets.
(2) M is inconsistent with respect to FB equivalence but consistent with respect to buffer equivalence (metrics with opposite consistency properties exist as well; see [8] for further details).

The recovery of such metrics via the argument described in Eq. (2) motivates the problem we study in this note: inverting the many-to-one map B and counting the size of its preimage. Results for map B are slightly stronger than those proven in [9] for map FB. Namely, computing the cardinality of the preimage of map B, as well as returning one element from the preimage can be done in polynomial time (even linear time for the latter problem).

3. Preliminaries

We will use notation $x \stackrel{\cdot}{-} y = \max\{x - y, 0\}$.

We employ standard graph theoretic notions throughout. In this paper, graphs are always bipartite and undirected.

Definition 5. A bipartite graph $G = (V_1, V_2, E)$ is *doubly convex* if there exist permutations π_1, π_2 of vertex sets V_1, V_2, respectively, such that for every $i \in \{1, 2\}$ and every vertex $v \in V_i$ the set of vertices w that are adjacent to v forms an interval (i.e. a set of consecutive nodes) of $\pi_{3-i}(V_{3-i})$.

Definition 6. A sequence of non-negative integers W is a *valid buffer pattern* if there exists a permutation A of $\{1, 2, \ldots, |W|\}$ such that $B(A) = W$.

Note that any valid buffer pattern W necessarily ends in a zero, since for $A \in B^{-1}(W)$ all packets in A can be passed to the application layer when the last packet in A is received. Also, without loss of generality, one can assume that *the only position in a valid buffer pattern that is equal to zero is the last one*, since one can decompose a given pattern W into disjoint segments, bounded by those positions equal to zero (where the buffer, therefore, gets flushed). With each such segment one can associate a permutation of a contiguous set of IDs.

Finally, rather than working with binary encodings of the input to our algorithm (in Fig. 1) for inverting buffer patterns, we will analyze its complexity in terms of the input length n of the list of numbers, and assume that operations on these numbers take $O(1)$ time. This is reasonable, since input numbers W_i are at most n in the case of a valid buffer pattern (this is an easy consequence of condition (5) of Theorem 4.1 (i)). Alternatively, one can could use binary encodings and still obtain a linear time algorithm if all input integers are represented using exactly $\lceil \log_2(n) \rceil$ bits.

Algorithm TwoStageGreedy(W)

INPUT: a list $W = (W_1, W_2, \ldots, W_n)$ of non-negative integers.

let σ be a vector of n numbers, initially all zero;
let ACK be a vector of $n + 1$ numbers, initially all zero
with the exception of $ACK_0 = 1$;
let *chosen* be an n-bit vector, with all positions initially zero;
let $W_0 = 0$;

for $i = 1$ to n
    if $(W_i - W_{i-1} > 1) \vee ((i < n) \wedge (W_i = 0)) \vee ((i = n) \wedge (W_i \neq 0))$
        reject
    else
      if $(W_i = W_{i-1})$
        let $\sigma[i] = ACK_{i-1}$;
        let $chosen[ACK_{i-1}] = 1$;
        let $ACK_i = ACK_{i-1} + 1$;
      else
        if $(W_i < W_{i-1})$
          let $\sigma[i] = ACK_{i-1}$;
          let $chosen[ACK_{i-1}] = 1$;
          let $ACK_i = ACK_{i-1} + W_i - W_{i-1} + 1$;
        else
          /* $W_i = W_{i-1} + 1$ */
          let $ACK_i = ACK_{i-1}$;

for $i = 1$ to n
    if $(\sigma[i] = 0)$
        let $\sigma[i] =$ the first $j > ACK_{i-1}$ with $chosen[j] = 0$;
        let $chosen[\sigma[i]] = 1$;

return σ.

Fig. 1. Algorithm for inverting buffer sequences.

4. Inverting buffer sequences

Our main result is

Theorem 4.1. *The following are true:*

(1) *A sequence of non-negative integers $W = (W_0, W_1, \ldots, W_n)$ is a valid buffer sequence* if and only if

$$W_0 = W_n = 0 \quad and \quad W_i \neq 0 \quad for\ all\ 1 \leq i < n, \quad and \tag{4}$$
$$W_i \leq W_{i-1} + 1\ for\ all\ 1 \leq i < n. \tag{5}$$

(2) *There is an algorithm that, given a sequence of non-negative integers $W = (W_1, W_2, \ldots, W_n)$ satisfying the conditions of Eqs. (4) and (5), constructs in time $O(|W|)$ a permutation A such that $A = B^{-1}(W)$.*

(3) *Counting the cardinality of the set of permutations in the preimage $B^{-1}(W)$ can be done in polynomial time.*

Proof. We will provide, in essence, a reduction of the problem above to the problem of finding a maximum matching in a special class of doubly convex bipartite graphs [12]. The complexity of this problem is linear in the number of vertices of the graph [12]. Since the size of the bipartite graph that is created by reduction is linear, the overall complexity of the problem is linear.

A valid buffer sequence consists of positive integers, with the exception of the last entry, which is zero. Any two consecutive values of the buffer sequence W_i and W_{i+1} can only be in one of the following situations:

(1) $W_i = W_{i-1} + 1$. This situation corresponds to one new out-of-order packet being received at stage *i*. This holds for $i = 1$ as well, if we let $W_0 = 0$.

(2) $W_i < W_{i-1}$. This situation corresponds to the newly received packet causing a non-empty portion of the buffer to be flushed. In particular the ID of the received packet can be inferred at this stage, and is equal to the smallest index of a packet not received so far.

(3) $W_i = W_{i-1}$. This situation corresponds to the packet received at this stage being the first packet not previously received. Receiving this packet does not cause any other packet to be sent to the application layer.

If the input sequence fails to satisfy these conditions (for instance if there exists an index i with $W_i - W_{i-1} > 1$) then the set of permutations in $B^{-1}(W)$ is empty. Otherwise, let S_1, S_2, S_3 be the set of indices corresponding to the three cases listed above.

During the course of the algorithm we will keep track of the value ACK_i, computed assuming that W is a valid buffer pattern. Initially $\text{ACK}_0 = 1$. We have the following recurrence relations (mirroring the three cases described above):

(1) The newly received packet is out of order. Thus, it does not change the value of parameter ACK. Therefore

$$\text{ACK}_i = \text{ACK}_{i-1}. \tag{6}$$

(2) The newly received packet has ID ACK_{i-1}. In addition, it makes the buffer shrink in size from W_{i-1} to W_i, which means that

$$\text{ACK}_i = \text{ACK}_{i-1} + W_{i-1} - W_i + 1. \tag{7}$$

(3) The newly received packet has index ACK_{i-1} and does not cause the buffer to shrink any more. Therefore

$$\text{ACK}_i = \text{ACK}_{i-1} + 1. \tag{8}$$

For all indices $i \in S_2 \cup S_3$, the index of the received packet is uniquely determined, and equal to ACK_{i-1}.

We will now create a bipartite graph $G = (V_1, V_2, E)$. Nodes in V_1 correspond to stage indices $i \in \{1, \ldots, n\}$. Nodes in V_2 will correspond to packet IDs. First, let $V_1 = S_1$, and let $V_2 = \{1, \ldots, n\} \setminus \{\text{ACK}_{i-1} | i \in S_2 \cup S_3\}$. Clearly $|V_2| = n - |S_2 \cup S_3| = |S_1| = |V_1|$. Second, given node $i \in V_1$, add edges to all vertices $j \in V_2$ such that $j > \text{ACK}_i$.

With this definition we have:

Lemma 4.1. *Permutations from the set $B^{-1}(W)$ are in bijective correspondence with elements of MATCH(G), the set of all perfect matchings in G. In particular $B^{-1}(W) \neq \emptyset$ if and only if G has a perfect matching.*

Proof. Each permutation can be seen as a set of pairs (i, j). By the previous discussion, the set of acknowledgements $\{\text{ACK}_i\}_{i \geq 0}$ is the same for any permutation in $B^{-1}(W)$. Moreover, for all $\sigma \in B^{-1}(W)$ and index $i \in S_2 \cup S_3$, $\sigma[i] = \text{ACK}_{i-1}$. Also, for such a permutation σ, by definition of graph G it is easy to see that all pairs $(i, \sigma[i])$ with $i \in S_1$ are edges in G. Hence σ corresponds to a perfect matching in G.

Conversely, every perfect matching M in G naturally corresponds to a sequence of pairs, that can be completed (by adding all pairs (i, ACK_{i-1}) for all values i not in V_1) to a mapping A defined on $\{1, \ldots, n\}$. A is actually a permutation. Indeed, the values of parameter ACK_i, $i \in S_2 \cup S_3$, are all different, and are not included in V_2. It follows that A maps n numbers onto n different numbers; hence it is a bijection.

To show that $A \in B^{-1}(W)$, assume that this was not the case, and let i be the smallest index such that $B_{A,i} \neq W_i$. Thus $B_{A,i-1} = W_{i-1}$ where, by convention, $B_{A,0} = 0$.

Case 1. $B_{A,i} = B_{A,i-1} + 1$. Since $W_i \neq B_{A,i}$ and $W_i - W_{i-1} \leq 1$, the only possible alternatives are $W_i = W_{i-1}$ or $W_i < W_{i-1}$. But then index i is not in V_1 and is matched in A to integer ACK_{i-1}. This contradicts the assumption that $B_{A,i} = B_{A,i-1} + 1$, since the packet with ID ACK_{i-1} is the first not received in the first $i - 1$ phases, and can thus be uploaded at stage i. The contradiction comes from our assumption that sequences $B(A)$ and W are different.

Similar arguments can be applied in the two remaining cases for the evolution of sequence $B(A)$, and the conclusion of the argument is that $A \in B^{-1}(W)$. □

Lemma 4.2. *Let $a_1 \geq a_2 \geq \cdots \geq a_m$ be the number of ones on the first, second, ..., mth row of M_G, the adjacency matrix of G (call $(a_1, \ldots, a_m)$ the type of M_G). Then we have*

(1) *G has a perfect matching if and only if for all $i = 1, \ldots, m$, $a_i \geq m + 1 - i$. When this condition holds, a perfect matching in G can be constructed by taking elements on the diagonal of M_G.*

(2) *The number of matchings in G is given by*

$$|\text{MATCH}(G)| = a_m(a_{m-1} \dot{-} 1)(a_{m-2} \dot{-} 2) \cdot \ldots \cdot (a_1 \dot{-} (m-1)). \tag{9}$$

Proof. Denote the cardinality of set MATCH(G) by $\Gamma(a_1, \ldots, a_m)$ (to highlight its dependency on parameters $a_1, \ldots, a_m$). Expand the permanent across the last row. Since $a_1, \ldots, a_{m-1}$ are all greater or equal to a_m, it follows that $\Gamma(a_1, \ldots, a_m)$ is the sum of the permanent of a_m minors, all of them of type $(a_1 \dot{-} 1, \ldots, a_{m-1} \dot{-} 1)$. Thus, $\Gamma(a_1, \ldots, a_m) = a_m \cdot \Gamma(a_1 \dot{-} 1, \ldots, a_{m-1} \dot{-} 1)$, and formula (9) immediately follows on noting that, for all $i \geq 1$, $(a \dot{-} (i-1)) \dot{-} 1 = a \dot{-} i$. □

We now complete the proof of Theorem 4.1.

(1) The necessity of the condition follows from the analysis on the dynamics of the ACK parameter at the beginning of the proof. Suppose, therefore, that conditions (4) and (5) in the statement of Theorem 4.1 (1) hold. Note that, from the definition of the ACK parameters, we have:

    (a) For all $1 \leq i \leq n$, $0 \leq \text{ACK}_i \leq i + 1$ (one cannot have received more than the first i packets in the first i stages).

    (b) If $\text{ACK}_i = i + 1$ then $W_i = 0$.

The second condition implies the fact that for a valid buffer pattern $\text{ACK}_i \leq i$ for all i. Since, for $i \in S_2 \cup S_3$, we match elements i and $\text{ACK}_{i-1} < i$, the following simple claim holds:

Claim 1. *Let $V_1 = \{i_1 < i_2 < \cdots i_m\}$ and $V_2 = \{j_1 < j_2 < \cdots < j_m\}$ be the elements of V_1 and V_2 in the definition of graph G. Then, for all $1 \leq k \leq m$,*

$$i_k \leq j_k. \tag{10}$$

Now, for all $i \in S_1$ we have $\text{ACK}_i + 1 = \text{ACK}_{i-1} + 1 \leq i$. Thus $\text{ACK}_{i_k} + 1 \leq i_k \leq j_k$, which means that the edge (i_k, j_k) is in $E(G)$. We obtain a perfect matching by taking the set of pairs (i_k, j_k).

(2) Algorithm TwoStageGreedy in Fig. 1 produces a perfect matching (if it exists). Its correctness follows from:

    (a) The previous discussion on the recurrence relations for the ACK parameters in the case of a valid pattern.

    (b) Point (1) of the Theorem.

    (c) Point (1) of Lemma 4.2.

    With a little care the algorithm can be implemented in $O(n)$ time. The only issue is finding the element j such that $\sigma[i]$ is set to j in the second **for** loop. Rather than starting from ACK_{i-1} each time, we use the observation that the matching can be computed as the diagonal of matrix M_G. That is, for every i with $\sigma[i] = 0$, $\sigma[i]$ can be taken the *smallest j* with $chosen[j] = 0$. The search for the next element will run from position j, and can be completed in time $O(n)$ by scanning elements of vector *chosen* from left to right.

(3) Computing $|\text{MATCH}(G)|$ using formula (9) can done in polynomial time as follows:

    (a) First, there is a polynomial time algorithm that, given input W, outputs the list of numbers $a_1, \ldots, a_m$. Indeed, the algorithm TwoStageGreedy computes parameters $\text{ACK}_1, \ldots, \text{ACK}_n$ in linear time.

    (b) Given these numbers, computing $|\text{MATCH}(G)|$ can be accomplished in time polynomial in $m + \lceil \log_2 |\text{MATCH}(G)| \rceil$ by the brute-force product computation in (9). Since $|\text{MATCH}(G)| \leq n!$ (simply because matchings correspond to permutations), it follows by Stirling's approximation that $\lceil \log_2 |\text{MATCH}(G)| \rceil = O(n \log n)$. Thus, the running time is polynomial in n. $\quad\square$

The proof of Theorem 4.1 also implies that buffer equivalence is a refinement of behavioral equivalence for permutations:

Proposition 1. *Let P and Q be two* permutations *such that $P \equiv_{buf} Q$. Then $P \equiv_{beh} Q$.*

Proof. Eqs. (6)–(8) show that the value of parameter ACK_i can be recovered directly from the buffer sizes. Since P and Q are buffer equivalent, they have identical buffer size sequences and, consequently, identical sequences of parameter ACK_i. But it is easy to see that the sequence of packet IDs (more precisely the corresponding sequence of byte IDs) ACKed by the TCP protocol in the case of simple consecutive ACKs is precisely ACK_i. Therefore P and Q are behaviorally equivalent. $\quad\square$

5. Reconstructing packet sequences with repeats

Buffer equivalence is not a refinement of behavioral equivalence in the presence of repeats. The reason is that one cannot distinguish between the case when the newly received packet is a repeat and Case 3 in the proof of Theorem 4.1 (in both cases the buffer size stays the same). However, for a repeat packet the value of the ACK parameter does not change, while for a packet in Case 3 the value of the ACK parameter increases by one.

We can modify (in a somewhat artificial way) the notion of buffer equivalence to incorporate information on whether the newly received packet is a repeat packet or not. For instance, one can define $B_{A,i}$ to be *minus* the buffer size when the ith received packet is a repeat. Denote this new mapping by $\bar{B}$.

Definition 7. Two sequences of packets P and Q are *modified buffer equivalent* (written $P \equiv_{\overline{buf}} Q$) if $\bar{B}(P) = \bar{B}(Q)$.

Similarly, we can modify the definition of valid buffer patterns. The analog of Theorem 4.1 for mapping $\bar{B}$ is

Theorem 5.1. *Let $W = (W_1, W_2, \ldots, W_n)$ be a sequence of non-negative integers.*

Deciding whether W is a valid buffer pattern, and in this case constructing an ID sequence A such that $A = \bar{B}^{-1}(W)$, can be done in time $O(n)$. Counting the cardinality of the preimage $\bar{B}^{-1}(W)$ can be done in polynomial time.

Proof. We only outline the proof, since it is very similar to that of Theorem 4.1. Given our use of negative numbers in the encoding, we no longer have the positivity constraint for elements of the candidate sequence W. However, we still require that only the last element be zero.

The construction of graph G is identical to that of Theorem 4.1, since in all stages in V_1 we can guarantee that a new packet is received. However, we do *not* have a parsimonious reduction of ID sequences to perfect matchings, since repeat packets can complete a matching in G in more than one way.

A polynomial time counting algorithm exists, nevertheless, since we can complement Lemma 4.2 with

Lemma 5.1. *We have*

$$|\overline{B}^{-1}(W)| = |MATCH(G)| \times \left(\prod_{i \in R} |W_i| \right), \tag{11}$$

where MATCH(G) is the set of all perfect matchings in G, and $R = \{i | W_i < 0\}$, i.e. the set of stages in which a repeat packet arrives. In particular $\overline{B}^{-1}(W) \neq \emptyset$ if and only if G has a perfect matching.

Also, the construction shows that modified buffer equivalence *is* a refinement of behavioral equivalence. Indeed, from the sequence of modified buffer sizes one can uniquely reconstruct the sequence of acknowledgments. The proof then proceeds just as the proof of Proposition 1. □

Acknowledgments

References

[1] G. Istrate, A. Hansson, S. Thulasidasan, M. Marathe, C. Barrett, Semantic compression of TCP traces, in: F. Boavida (Ed.), Proceedings of the Fifth IFIP NETWORKING Conference, in: Lecture Notes in Computer Science, vol. 3976, Springer-Verlag, 2006, pp. 123–135.

[2] S. Jaiswal, G. Iannacone, C. Diot, J. Kurose, D. Towsley, Inferring TCP connection characteristics through passive measurements, in: Proc. IEEE INFOCOM Conference, Hong Kong, China, 2004.

[3] J.C.R. Bennett, C. Partridge, N. Shectman, Packet reordering is not pathological network behavior, IEEE/ACM Transactions on Networking 7 (6) (1999) 789–798.

[4] W. Stevens, TCP/IP Illustrated, Vol.1: The Protocols, Addison-Wesley, 1994.

[5] T. Banka, A.A. Bare, A.P. Jayasumana, Metrics for degree of reordering in packet sequences, in: Proc. 27th IEEE Conference on Local Computer Networks, 2002, pp. 333–342.

[6] N.M. Piratla, A.P. Jayasumana, A.A. Bare, Reorder Density (RD): A formal, comprehensive metric for packet reordering, in: Raouf Boutaba, et al. (Eds.), Proceedings of the Fourth IFIP Networking Conference 2005, in: Lecture Notes in Computer Science, vol. 3462, Springer-Verlag, 2005, pp. 78–89.

[7] A. Morton, L. Ciavattone, G. Ramachandran, S. Shalunov, J. Perser, Packet reordering metrics. IETF IPPM Working Group RFC 4737, available from http://www.ietf.org/rfc/rfc4737.txt (Last accessed March 2008).

[8] A. Hansson, G. Istrate, G. Yan, Packet reordering metrics: Some methodological considerations, in: Proceedings of the Second International Conference on Networking and Services (ICNS'06), I.E.E.E. Computer Society Press, ISBN: 0-7695-2622-5, 2006.

[9] A. Hansson, G. Istrate, S. Kasiviswanathan, Combinatorics of TCP reordering, Journal of Combinatorial Optimization 12 (1–2) (2006) 57–70.

[10] J. Kurose, K. Ross, Computer Networking: A Top-Down Approach Featuring the Internet, second ed., Addison-Wesley, 2003.

[11] A.P. Jayasumana, N.M. Piratla, A.A. Bare, T. Banka, R. Whitner, J. McCollom, Reorder density and reorder buffer-occupancy density - metrics for packet reordering measurements, IETF draft, available from http://www.ietf.org/internet-drafts/draft-jayasumana-reorder-density-08.txt (Last accessed March 2008).

[12] W. L. Jr., F. Preparata, Efficient algorithms for finding maximum matchings in convex bipartite graphs and related problems, Acta Informatica 15 (1981) 329–346.