

COMMENT ON “A DETERMINISTIC APPROACH TO THE END-TO-END ANALYSIS OF PACKET FLOWS IN CONNECTION ORIENTED NETWORKS”

JEAN-YVES LE BOUDEC
GÉRARD HÉBUTERNE

Abstract

We prove that the buffer bound in [CFZF98] can be improved, by using a modification of the proofs in the original paper [CFZF98], together with so-called network calculus bounds. We also show that the delay bound in [CFZF98] is the sum of worst case queuing delays at all nodes along the path of a connection.

Contents

1	Introduction	1
2	Improving the bounds in [CFZF98]	2
3	Concluding Remarks	5
A	Appendix: Proof of Theorem 2.2	6

1 Introduction

In [CFZF98], the authors consider a network of discrete time, First In First Out (FIFO) queues. They assume that the network uses a connection oriented paradigm, and that packets (called “cells”) all have the same size (as is the case with ATM). In particular, it is assumed that all cells belonging to one connection follow the same path, established at connection setup. In this context, the words “connection” and “flow” have the same meaning. It is further assumed that a connection is spaced at the source by at least the Route Interference Number (RIN) of the connection. The RIN of connection r is defined as the number of occurrences of other connections joining the path of r . This assumption is called the *source rate condition*; it is shown in [CFZF98] that under this set of assumptions:

- the network is stable (namely, queue lengths remain bounded);

- the delay for any connection is bounded by its RIN;
- the maximum buffer required at a queue with I input links and N_i connections on input link i is bounded by $\max_{1 \leq i \leq I} (N - N_i)$, with $N = \sum_{i=1}^I N_i$.

In this contribution, we show that, under the same assumptions, it is possible to improve the buffer bound to $\min_{1 \leq i \leq I} (N - N_i)$ instead of $\max_{1 \leq i \leq I} (N - N_i)$. This can be achieved using a small variation of network calculus bounds [Cru95, Bou98, Cha98, RR96], together with a modification of the main proofs in [CFZF98]. This also implies some improvements for the delay bound.

Essentially the same result was found independently by H. Zhang in [Zhn99], using a different approach, based on a detailed analysis of worst case delays. See also some concluding remarks in Section 3. In [Zhn99], Zhang also analyzes the tightness of the bound.

In Section 2 we give the new bounds. Section 2 rely on a number of lemmas, which are given in appendix.

Throughout the paper we adopt a discrete time model, as in [CFZF98], and assume that all packets have a the same size, equal to one unit of data.

2 Improving the bounds in [CFZF98]

Our starting point is a number of results which we collectively refer to as “network calculus” [Cru95, Bou98, Cha98, RR96]. These results give deterministic bounds on buffer and delay, assuming input processes are limited by some arrival curves, and the service element offer some form of service guarantee. We say that a flow admits a function $\alpha(t)$ as arrival curve if the number of cells that can be observed on the flow during *any* interval of duration t is $\leq \alpha(t)$. We also say that a node e offers to a flow r a “strict service curve” $\beta(t)$ if, during any time interval of length t , for which the backlog of connection r at node e is positive, the number of cells of flow r which are output by the node is $\geq \beta(t)$. The backlog for flow r at node e at some time instant is defined as the number of cells of flow r which have entered node e and did not depart yet. The strict service curve property was defined for example in [LM96] and is an abstraction of the generalized processor sharing concepts introduced in [KG94]. The following theorem is a new variant of classical results in [Cru95, Bou98, Cha98, RR96].

Theorem 2.1. *Consider a node which receives an input connection, with a buffer large enough to avoid discarding data. Assume that the node offers a strict service curve β and that the input connection has an arrival curve α . Assume that $\alpha(u_0) \leq \beta(u_0)$ for some $u_0 > 0$. Then the maximum buffer occupancy is bounded by $\sup_{0 \leq u < u_0} (\alpha(u) - \beta(u))$*

The theorem says that, for the computation of a buffer bound, it is sufficient to consider time intervals less than u_0 . The idea is that the busy period duration is less than u_0 .

Proof: The proof is similar to network calculus bounds in [Cru95, Bou98, Cha98, RR96]. Call $x(t)$ [resp. $x^*(t)$] the input [resp. output] function. This is traditionally defined as the cumulative number of cells observed on the input [resp. output] flow over the time interval $[0, t]$. Consider a given time t at which the buffer is not empty, and call s the last time instant before t at which the buffer was empty. Then, from the strict service curve property, we have

$$x^*(t) \geq x^*(s) + \beta(t - s) = x(s) + \beta(t - s)$$

Thus the buffer size at time t satisfies

$$b(t) = x(t) - x^*(t) \leq x(t) - x(s) - \beta(t - s) \leq \alpha(t - s) - \beta(t - s)$$

Now if $t - s \geq u_0$, then there exists a time $t' = s + u_0$, with $s + 1 \leq t' \leq t$ such that $b(t') = 0$. This contradicts the definition of s . Thus we can assume that $t - s < u_0$. \square

Now we proceed with a property which generalizes the results in [CFZF98], and which will be required for improving the buffer bound.

Theorem 2.2. *With the assumptions in [CFZF98], consider a given link e and a subset \mathcal{S} of m connections which use that link. Let n be a lower bound on the number of route interferences that any connection in the subset will encounter after this link. Then over any time interval of duration $m + n$, the number of cells belonging to \mathcal{S} that leave link e is bounded by m .*

An equivalent way to formulate the theorem is to say that, if we call α the minimum arrival curve for the aggregate of the m flows on link e , then we have

$$\alpha(m + n) \leq m \tag{1}$$

Before giving the proof, let us mention the following. For $m = 1$, the theorem means that the spacing between cell departures from link e is at least $n + 1$, where n is the "remaining" route interference number for the connection, past link e . In general, n is less than the RIN of the connection. The difference $n - RIN$ is due to cell delay variation accumulated in the buffers along the path. In other words, a connection gains cell delay variation along its path, and the cell delay variation is bounded by the route interference number consumed along the path. This result (namely for $m = 1$) derives immediately from Theorem 1 in [CFZF98]

However the theorem is more powerful. It gives a bound on the number of arrivals for an aggregate flow. It can easily be seen that the bound is not a simple consequence of bounds for individual connections; indeed, the bound so obtained would be $\alpha(1 + n) \leq m$, instead of $\alpha(m + n) \leq m$. In contrast, the bound in the theorem takes into account global interactions between connections.

Proof: Remember that we have assumed a connection oriented network; thus every cell (= packet) belongs to one connection, and every connection uses one path. Consider an arbitrary cell c , and call r the connection it belongs to; as a shortcut, we also say that cell c "is in \mathcal{S} " if r is in \mathcal{S} .

Consider now a fixed time interval $(s, t] = [s + 1, t]$ with $t = s + m + n$. Call \mathcal{A} the set of all cells in \mathcal{S} which leave the link during $(s, t]$. Note that connections in \mathcal{S} may interfere at several different links, but since they all end up using link e , there is always one last link before or at e at which they interfere. We call this link the merging point of the two connections. We use the classical definition of busy period used in queuing theory, namely, a time interval during which the backlog for the flow at the node is always positive. For two cells c and d in \mathcal{S} , and for some link f , we say that $d \preceq c$ if c and d are in the same busy period at the merging points for the connections of c and d (see Figure 1). We will use the binary relation \preceq as follows. By Lemma 1 in [CFZF98], the delay for a cell c in \mathcal{S} due to interferences in \mathcal{S} at the merging point is bounded by the number of cells d in \mathcal{S} satisfying $d \preceq c$.

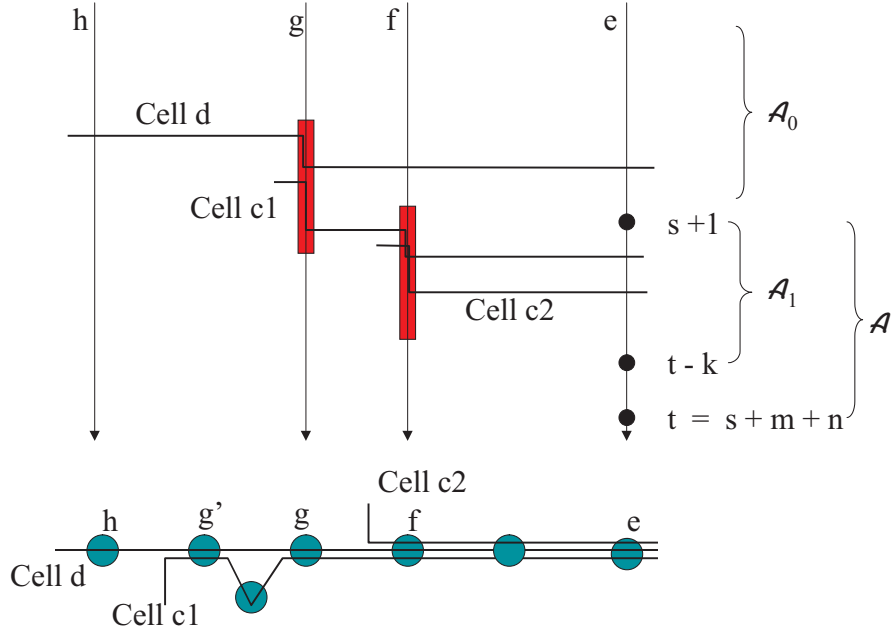


Figure 1: A time-space diagram illustrating the definitions of $d \preceq c_1$ and $c_1 \preceq c_2$ (all cells shown are in \mathcal{S}). Time flows downwards. Rectangles illustrate busy periods. c_1, c_2 are in \mathcal{A}_1 while d is in \mathcal{A}_0 . The merging point for the flows of d and c_1 is link g (not g').

Our definition of $d \preceq c$ is very close to the concept of *delay chain* used in [CFZF98]. More precisely, $d \preceq c$ at some link f is equivalent to saying that cells c and d are in the same delay chain at link f , that link f is the merging point for the connections of c and d , and that d reached link f before c .

We now define the set \mathcal{A} as follows; we say that some cell d is in \mathcal{A}_0 if and only if d is in \mathcal{S} , d leaves link e before or at time s , and there exists a sequence of cells $c_0 = d, c_1, \dots, c_k$ all in \mathcal{S} , such that $c_{i-1} \preceq c_i$ for $i = 1, \dots, k$, and $c_k \in \mathcal{A}$. The definition of \mathcal{A}_0 is similar to that of the *superchain* for a given path in [CFZF98]; however it differs in that it does not apply to one specific path, and that the cells in the chain are restricted to be in \mathcal{S} .

Call $k = \#\mathcal{A}_0$, where the $\#$ sign indicates the number of elements in a set (thus k is the number of connections that are in \mathcal{A}_0). Note that it is possible that \mathcal{A}_0 is empty, in which case $k = 0$. Also call \mathcal{A}_1 the subset of \mathcal{A} made of those cells that leave the link in $(s, t - k]$. Since the link rate is one cell per time unit, we have:

$$(\#\mathcal{A}) \leq k + (\#\mathcal{A}_1) \quad (2)$$

The main idea of the proof is that

$$(\#\mathcal{A}_1) \leq m - k \quad (3)$$

which together with Equation (2) will prove the result. Equation (3) follows from Lemma A.1 which shows that there can be at most one cell per connection in $\mathcal{A}_0 \cup \mathcal{A}_1$. \square

We now come to our main result. Consistent with the network model in [CFZF98], a network nodes is modelled as a collection of output buffers, with no contention other than at the output buffers. Every buffer is associated with one unidirectional link which

it feeds. Every link has one origin node and one end node. We say that a link f is incident to link e if the origin node of link e is the destination node of link f . In general, a link has several incident links.

Theorem 2.3. *Consider some arbitrary link e with I incident links, and call N_i the number of connections that use link e and arrive on the i th incident link, $i \leq I$. Also call $N = \sum_{i=1}^I N_i$. With the assumptions in [CFZF98], the amount of data in the buffer is bounded by $\min_{1 \leq i \leq I} (N - N_i)$ (instead of $\max_{1 \leq i \leq I} (N - N_i)$ as in [CFZF98])*

Proof: We apply Theorem 2.2 to the set of connections that arrive at e on link i , with $m = N_i$ and $n = N$. The maximum number of cells that can arrive during t time slots at link e via the i th incident link is thus limited by $\alpha_i(t) = \min(t, N_i)$, for $t \leq N$. We now apply Theorem 2.1 to $\alpha = \sum_{i=1}^I \alpha_i$ and the strict service curve $\beta(t) = t$. We can assume without loss of generality that $N_1 \leq N_2 \leq \dots \leq N_I$. The function $\alpha - \beta$ is continuous and has a derivative at all points except the N_i 's. The derivative changes its sign at N_I ($= \max_{1 \leq i \leq I} (N_i)$) thus the maximum of $\alpha - \beta$ is at N_I and its value is $N - N_I$. \square

The delay bound Lastly, let us discuss the delay bound. Call $j(r, e)$ the link by which connection r arrives at node e . From Theorem 2.3, the delay experienced by a cell of connection r arriving at node e is bounded by

$$\min_{1 \leq i \leq I(e)} (N(e) - N_i(e)) \leq N(e) - N_{j(r,e)}(e) \quad (4)$$

Here we have denoted with $I(e)$ the number of incident links at node e , $N_i(e)$ the number of connections arriving at node e on link i , and $N(e) = \sum_{i=1}^{I(e)} N_i(e)$. Now $N(e) - N_j(e)$ is the number of route interferences for connection r at node e . Also write $e \in r$ to express that node e is on the path of connection r . The end-to-end delay for connection r is thus bounded by

$$\delta(r) = \sum_{e \in r} (N(e) - N_{j(r,e)}(e)) \quad (5)$$

which is precisely the RIN of connection r .

This result is already in the original paper [CFZF98]. However, we should mention here firstly that, contrary to what might be interpreted from [CFZF98], the end-to-end delay bound is the sum of the local, independent delay bounds at every node. Secondly, a better bound can be directly obtained by using the left-hand side in Equation (4) instead of the right-hand side. This gives the following bound for the end-to-end delay:

$$\delta'(r) = \sum_{e \text{ such that } e \in r} \left\{ \min_{i \text{ such that } 1 \leq i \leq I(e)} (N(e) - N_i(e)) \right\}$$

Namely, the end-to-end delay is bounded by the sum of the minimum numbers of route interferences for all connections at all nodes along the path of a connection. For asymmetric cases, this is less than the RIN of the connection as given in Equation (5).

3 Concluding Remarks

Essentially the same result was found independently by H. Zhang in [Zhn99], using a different approach, not based on network calculus. In our approach, we show an intermediate result (Theorem 2.2), which gives a property of the arrival function for an aggregate

number of connections. We believe that this direction could be used to analyze generalizations of the original problem in [CFZF98], in particular, if we consider more general general source rate conditions.

References

- [Bou98] J.-Y. Le Boudec. Application of network calculus to guaranteed service networks. *IEEE Transactions on Information Theory*, 44:1087–1096, May 1998.
- [CFZF98] I. Chlamtac, A. Faragó, H. Zhang, and A. Fumagalli. A deterministic approach to the end-to-end analysis of packet flows in connection oriented networks. *IEEE/ACM transactions on networking*, (6)4:422–431, 08 1998.
- [Cha98] C.S. Chang. On deterministic traffic regulation and service guarantee: A systematic approach by filtering. *IEEE Transactions on Information Theory*, 44:1096–1107, August 1998.
- [Cru95] R.L. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE JSAC*, pages 1048–1056, August 1995.
- [KG94] Parekh A. K. and Gallager R. G. A generalized processor sharing approach to flow control in integrated services networks: The multiple node case. *IEEE/ACM Trans. Networking*, vol 2-2, pages 137–150, April 1994.
- [LM96] Cruz R. L. and Okino C. M. Service guarantees for window flow control. In *34th Allerton Conf of Comm., Cont., and Comp. Monticello, IL*, Oct 1996.
- [RR96] Agrawal R. and Rajan R. Performance bounds for guaranteed and adaptive services, December 1996. IBM Technical Report RC 20649.
- [Zhn99] Hongbiao Zhnag. A note on deterministic end-to-end delay analysis in connection oriented networks. In *Proc of IEEE ICC'99, Vancouver, pp 1223–1227*, 1999.

A Appendix: Proof of Theorem 2.2

We use a discrete time model, and assume that all propagation times are 0. The proof in this appendix can easily be modified to incorporate propagation times, but we prefer to leave this to the reader as this complicates a notation which is already complex enough.

The main technical result is the following lemma. It is an extension of the “excluded superchain” lemma in [CFZF98].

Lemma A.1. *With the notation in Theorem 2.2, there is at most one cell per connection in $\mathcal{A}_0 \cup \mathcal{A}_1$.*

Proof: Firstly, we prove that there cannot be two cells of the same connection in \mathcal{A}_0 . We proceed by contradiction. Assume that $d, d' \in \mathcal{A}_0$ belong to the same connection and were emitted in that order. There exists $c_0 = d, c_1, \dots, c_k$ in \mathcal{S} such that $c_0 \preceq c_1, c_1 \preceq c_2, \dots, c_{k-1} \preceq c_k$, with $c_k \in \mathcal{A}_1$. Call f_i the merging point for the connections of c_{i-1} and c_i . We show now that cell d' must be flowing on the common subpath after cell c_k . Indeed, otherwise, from the FIFO property, it would have reached f_k before c_k . Assume it has reached f_k after c_{k-1} ; then it would belong to a super chain from d to d' . This is




Figure 2: Two cells of same connection cannot be in \mathcal{A}_0




Figure 3: There cannot be two cells of the same connection in \mathcal{A}_1 .

impossible from Lemma 2 in [CFZF98]. Thus d' reaches f_k after c_{k-1} . By recursion, this shows that d' must have reached f_1 before $c_0 = d$ which is a contradiction. Thus cell d' must be flowing on the common subpath after cell c_k . Now this contradicts the facts that $d' \in \mathcal{A}_0$ and $c_k \in \mathcal{A}_1$.

Secondly, we show that there cannot be two cells of the same connection in \mathcal{A}_1 . Let d be a cell in \mathcal{A}_1 . Call $s + \tau$, with $\tau \geq 1$ the time at which d leaves the link. From Lemma 1 in [CFZF98] the delay experienced by d along its path is bounded by the number of interference units experienced by d . Now let us make the distinction between an interference unit which is due to a connection in \mathcal{S} and which occurs at the merging point with the connection of d (call r the number of such interference units experienced by c), and other interference units experienced by d (call r' their number). For example on Figure 3, the interference of c_1 at g' is counted in r' , whereas the interferences of c_1 at g and c_2 at f are counted in r . The delay experienced by d between its source and link e is thus bounded by $r + r'$. Now if c is an interfering cell counted in r , we have $c \preceq d$. Thus either $c \in \mathcal{A}_0$ or $c \in \mathcal{A}_1$. In the latter case, by the FIFO property, c must leave e

before d . Thus

$$r \leq k + \tau - 1 \quad (6)$$

Also call R' the number of route interferences for the connection of d due to connections that are either not in \mathcal{S} or are in \mathcal{S} but are not at the merging point with the connection of d . We have thus $r' \leq R'$. Call t_0 the departure time for cell d at its source. From the above we have

$$t_0 = s + \tau - r - r' \geq s - k + 1 - R'$$

Call t_1 the emission time for the next cell, say d_1 , following d at its source. By the assumptions in the theorem, the total number of route interferences for the connection of d is at least $(m - 1) + R' + n$, thus

$$t_1 \geq t_0 + m + R' + n$$

combining the two previous equations we have:

$$t_1 \geq s + m + n - k + 1 = t + k + 1$$

Thus, from the definition of \mathcal{A}_1 , d_1 is not in \mathcal{A}_1 , which proves that there cannot be two cells of the same connection in \mathcal{A}_1 .

Thirdly, we prove that if $d \in \mathcal{A}_0$ and $d' \in \mathcal{A}_1$, then necessarily d and d' belong to different connections. Consider some $d \in \mathcal{A}_0$, and a sequence $c_0 = d \preceq c_1, c_1 \preceq c_2, \dots, c_{k-1} \preceq c_k$, with all cells in \mathcal{S} and $c_k \in \mathcal{A}_1$. Call f_i the merging point for c_{i-1} and c_i . See Figure 4. Call b the arrival time of d at f_1 and $s + \tau$, with $\tau \geq 1$ the departure time for c_k at link e . Define r_i as the number of interference units in the busy period at f_i where c_{i-1} and c_i interfere, plus those experienced by cell c_i between f_i and f_{i+1} , excluding f_{i+1} , due to cells in \mathcal{S} at their merging point with the connection of d . Following the same reasoning as in the proof of Lemma 2 in [CFZF98] we have:

$$s + \tau - b \leq \sum_{i=1}^k (r_i + r'_i) \quad (7)$$

Call $t(d), t(d')$ the departure times of cells d and d' at their (common) source. We have similarly:

$$b - t(d) \leq r_0 + r'_0 \quad (8)$$

where r_0 [resp. r'_0] is the number of interferences units for cell d on its path from the source to f_1 , excluding f_1 , due to cells in \mathcal{S} at their merging point with the connection of d [resp. due to other cells].

Now since the spacing between cells at the source is at least the route interference number for the path of d and d' , we have:

$$t(d') - t(d) \geq m + n + R' \quad (9)$$

where R' is the number of route interferences for the path of d , due to connections not in \mathcal{S} at the merging point with the connection of d . Thus

$$\sum_{i=0}^k r'_i \leq R' \quad (10)$$

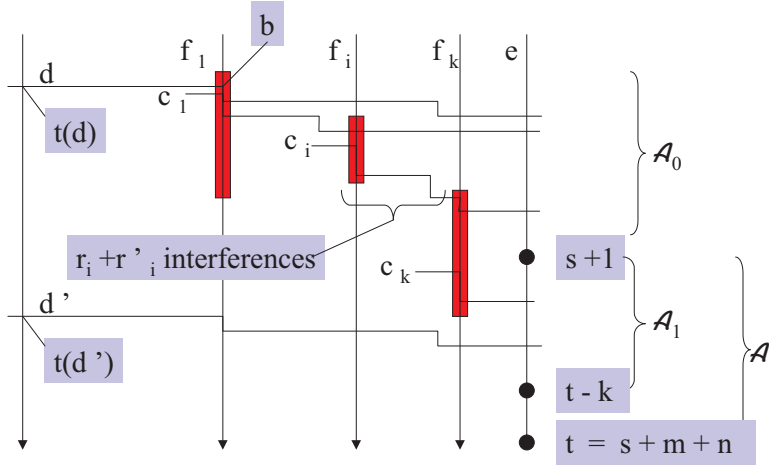


Figure 4: There cannot be two cells of the same connection in $\mathcal{A}_0 \cup \mathcal{A}_1$.

Combining Equations (7), (8), (9) and 10, we get:

$$t(d') \geq s + \tau + m + n - \sum_{i=0}^k r_i$$

Now every interference counted in r_i corresponds to one cell γ in \mathcal{S} with $\gamma \preceq c_i$ for some i . Thus either $\gamma \in \mathcal{A}_0$ or $\gamma \in \mathcal{A}$. By the FIFO property, in the latter case γ must leave link e before c_k . Thus $\sum_{i=0}^k r_i \leq k + \tau - 1$. Thus $t(d') \geq s + m + n - k + 1 = t - k + 1$ thus d' is not in \mathcal{A}_1 .

Combining the three arguments, we find that there is at most one cell per connection in $\mathcal{A}_0 \cup \mathcal{A}_1$. \square