Digital Fountain Technical Report DF2002-09-001

On WEBRC Wave Design and Sender Implementation

Vivek K Goyal
39141 Civic Center Drive, Suite 300
Fremont, CA 94538
vivek@digitalfountain.com

September 12, 2002

Abstract

A WEBRC sender transmits packets on a base channel and on `T` wave channels. The transmission pattern on each wave channel is a sequence of periodically-spaced waves separated by quiescent periods. In a fluid model, the rate of each wave is predominantly described by an exponential decay by a factor of `P` each `TSD` seconds. The beginnings of the waves are designed so that the sum over all of the channels of the fluid-model rate is constant. This design is complicated by requirements motivated by keeping receiver MRTT measurements simple.

Presented here is the design of the beginnings of the waves and a mechanism to convert fluid-model rates to packet transmission times. The final design has two parameters to trade off the range of fine-grained rate control against the maximum interpacket spacing in the early portion of the wave. For the sake of brevity, this document is far from self-contained. Those unfamiliar with WEBRC should see [1, 2] for an introduction.

1 Introduction

This document primarily addresses the design of a prototype wave that ends at time $\mathtt{N} \cdot \mathtt{TSD}$. We will say that the wave starts at time 0 whether or not the rate of wave is immediately greater than zero. (Note that this obvious choice of time coordinates places time slot boundaries at integer multiples of TSD.) With appropriate shifting by multiples of TSD, the prototype wave design results in a fluid-model transmission rate for every wave channel. The final part of this document addresses the scheduling of (discrete) packet transmissions to approximate the fluid model.

2 Fluid-Model Wave Design Problem

Many of the rate control mechanisms of WEBRC depend on the sum of channel rates being exponential with a multiplicative decay rate that is independent of the time within a time slot and the number of channels that the receiver is subscribed to. (To fix terms: the "multiplicative decay rate" is a factor of P per TSD seconds.) The fact that shifting in time does not change a multiplicative decay rate is what makes the exponential wave shape appropriate.

Most, if not all, of the wave has a simple exponential form:[1]

$$R(t) \;=\; \mathtt{P}^{-\mathtt{N}+t/\mathtt{TSD}} \cdot \mathtt{BCR}. \tag{1}$$

Define the *crest time* of the wave t_{crest} as the time after which the wave follows (1). The rate on the base channel follows the same exponential decay as the waves, but with a period of TSD. The base channel rate is given by

$$R_{\text{base}}(t) \;=\; \mathtt{P}^{\langle t/\mathtt{TSD}\rangle} \cdot \mathtt{BCR}, \tag{2}$$

where $\langle x \rangle$ is the fractional part of x; *i.e.*, $\langle x \rangle = x - \lfloor x \rfloor$.

The original design of WEBRC, as described in [3, 4], did not specify how waves should start or whether waves should start at time slot boundaries. Implicitly from the definition of MSR, one could conclude that waves would generally not start at time slot boundaries. Equivalently, a prototype wave with support $t \in [0, \mathtt{N} \cdot \mathtt{TSD}]$ has rate 0 for $t < t_{\text{crest}}$ where $t_{\text{crest}} \in [0, \mathtt{TSD})$. An example of the cumulative sender output with this original design is shown in Fig. 1.

Many figures similar to Fig. 1 appear in this document, so it is important to understand how to interpret them. The horizontal dimension represents the passage of time, and black vertical lines indicate time slot boundaries. Colors represent channel identities, where black is used for the base channel and all other colors represent wave channels. Heights represent fluid-model rates and the "fluid" for different channels is stacked so that being one higher in the stack corresponds to having a channel number one larger (modulo T). Since the fluid is stacked, the total height is the total instantaneous sending rate and the heights in different colors give the rates on the various active channels. In this figure, the prototype wave is tan because it is the wave that starts in the first time slot shown. Since the tan wave ends at time $10 \cdot \mathtt{TSD}$, one can conclude that $\mathtt{N} = 10$. The "earlier" or "lower" waves are brown, gray, magenta, ...; the "later" or "higher" waves are steel blue, royal blue, violet, ... Looking at any one color in Fig. 1, the wave channel rate appears at first to be discontinuous. This is an illusion caused by the stacking, the discontinuity in the base channel rate at time slot boundaries, and the discontinuous manner in which a wave ends at each time slot boundary.

The total sender output with this wave design is periodic with period TSD and may be far from constant; the ratio between the maximum and minimum rates can be as large as $1 + \mathtt{P}^{-1}$ (though

[1] All rates are in packets per second.

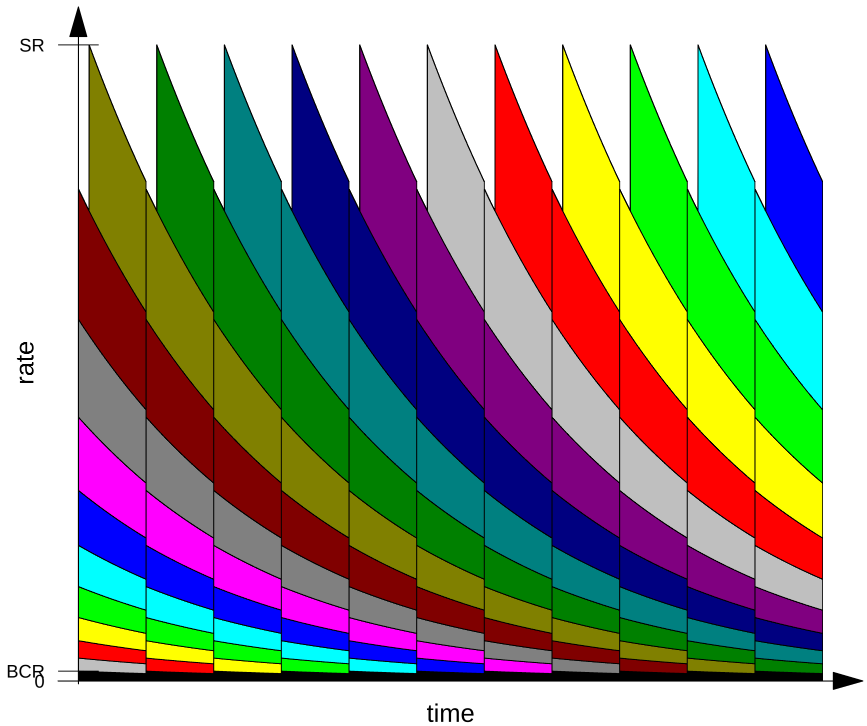

Figure 1: Example of cumulative output with the original design ($\mathtt{P} = \frac{3}{4}$, $\mathtt{SR}/\mathtt{BCR} = 65$).

$\mathtt{P}^{-1}$ is a more typical value). This periodicity has various disadvantages including incomplete usage of sender bandwidth and buffering and synchronization problems within the sender. This motivates filling in the gaps between saw teeth to make the cumulative output have constant rate SR.

More germane to the mechanisms of WEBRC, consider what happens when a receiver joins a highest wave early in the time slot in which it becomes active. Let time 0 be the beginning of the time slot at the network element where the receiver in question is grafted to the multicast tree.[2] Then let t_{join} be the time at which the receiver joins, t_{up} be the total propagation and processing times for all necessary IGMP and PIM messages, and t_{down} be the propagation time for data packets. The packet flow first arrives at the receiver at time $\max\{t_{\text{crest}}, t_{\text{join}} + t_{\text{up}}\} + t_{\text{down}}$. The normal mechanism for computing MRTT yields

$$\mathtt{MRTT} = [\max\{t_{\text{crest}}, t_{\text{join}} + t_{\text{up}}\} + t_{\text{down}}] - t_{\text{join}}$$

though the desired result is $t_{\text{up}} + t_{\text{down}}$. Thus MRTT may be overestimated by as much as TSD and the ratio between the measured and desired MRTTs can be arbitrarily large. This is unacceptable, so the wave design or MRTT calculation must be changed. Note that an accurate MRTT calculation by the receiver requires not only t_{crest}, which would probably be fixed throughout a session, but also an estimate of the join time. Thus far, the design of WEBRC has avoided tracking time within time slots as much as possible. For the purposes of this document, assume that the MRTT calculation is not changed.

The design objectives are to create a prototype wave that makes the total sender output constant while also reducing MRTT measurement errors.

[2] In the terminology of [2], the computations are using *wave time* at the network element where the receiver is grafted to the multicast tree.

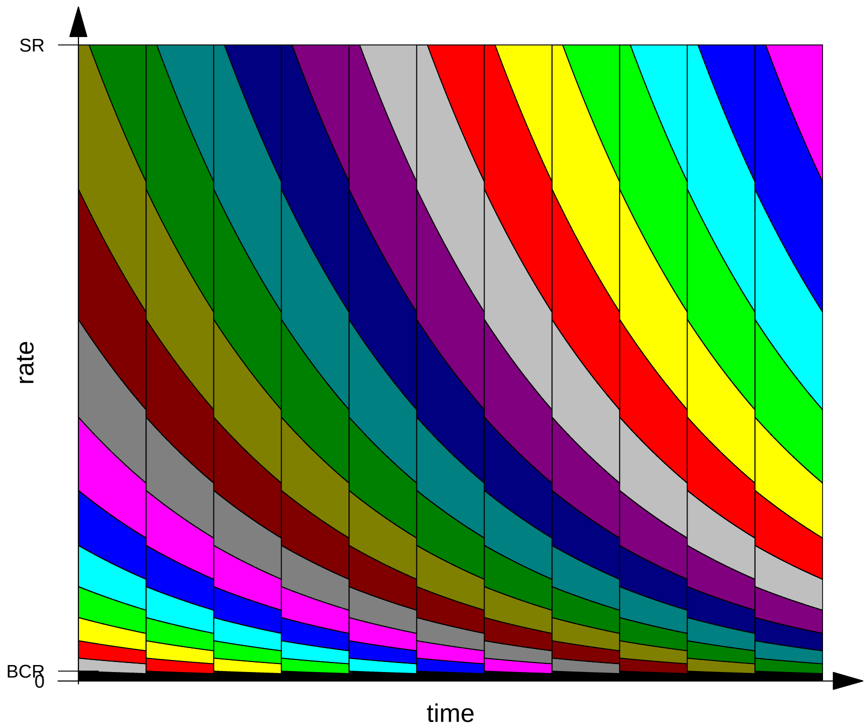

Figure 2: Example of Design 1 ($\mathtt{P} = \frac{3}{4}$, $\mathtt{SR}/\mathtt{BCR} = 65$).

3 Preliminary Designs

3.1 Design 1: Simple Extension

Perhaps the most obvious way to adjust the wave design from Fig. 1 is to extend waves to the left to fill the fluid level up to $\mathtt{SR}$. The result is shown in Fig. 2. Notice that this adds one time slot to the active duration of a wave. The resulting $\mathtt{N}$ is given by

$$\mathtt{N} \;=\; \left\lceil \log_{1/\mathtt{P}} \left(1 + \frac{1}{\mathtt{P}} \left(\frac{1}{\mathtt{P}} - 1 \right) \frac{\mathtt{SR}}{\mathtt{BCR}} \right) \right\rceil - 1, \tag{3}$$

as in [1, 2]. This value is derived from the following fact: $\mathtt{N}$ should be smallest integer such that the rate at the end of each time slot, without truncation of waves,

$$\mathtt{BCR} \left(\mathtt{P} + 1 + \frac{1}{\mathtt{P}} + \cdots + \frac{1}{\mathtt{P}^{\mathtt{N}-1}} \right),$$

is at least $\mathtt{SR}$.

By comparing Figs. 1 and 2, one can conclude that the prototype wave for Design 1 is given by

$$R(t) \;=\; \begin{cases} 0, & 0 \le t < t_{\text{crest}} - \mathtt{TSD}; \\ \mathtt{SR} - \dfrac{\mathtt{P}^{-\mathtt{N}} - 1}{\mathtt{P}^{-1} - 1} \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR}, & t_{\text{crest}} - \mathtt{TSD} \le t < \mathtt{TSD}; \\ \mathtt{SR} - \dfrac{\mathtt{P}^{-(\mathtt{N}-1)} - 1}{\mathtt{P}(\mathtt{P}^{-1} - 1)} \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR}, & \mathtt{TSD} \le t < t_{\text{crest}}; \\ \mathtt{P}^{-\mathtt{N}+t/\mathtt{TSD}} \cdot \mathtt{BCR}, & t_{\text{crest}} \le t \le \mathtt{N} \cdot \mathtt{TSD}. \end{cases} \tag{4}$$

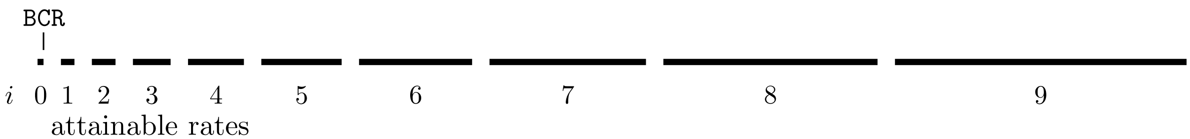

Figure 3: Rates that are attainable with only exponential portions of waves ($\mathtt{P} = \frac{3}{4}$). Segments are labeled with values of i as in (5). Segment lengths increase geometrically with fixed gaps of length $\mathtt{P} \cdot \mathtt{BCR}$.

The second of the four cases above is obtained by subtracting the sum of the rates of $\mathtt{N} - 1$ waves and the base channel from the total rate $\mathtt{SR}$. Specifically, for a time $t_{\text{crest}} - \mathtt{TSD} \leq t < \mathtt{TSD}$ the rate on the base channel and all lower waves is

$$\underbrace{\mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR}}_{\text{base channel}} \quad + \quad \underbrace{\mathtt{P}^{-1} \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR} + \mathtt{P}^{-2} \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR} + \cdots + \mathtt{P}^{-(\mathtt{N}-1)} \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR}}_{\text{wave channels}}$$
$$= \left(\sum_{k=0}^{\mathtt{N}-1} \mathtt{P}^{-k} \right) \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR} \;=\; \frac{\mathtt{P}^{-\mathtt{N}} - 1}{\mathtt{P}^{-1} - 1} \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR}.$$

The third case in (4) is obtained similarly by adding the rates of $\mathtt{N} - 2$ wave channels and the base channel.

One subtlety in this description—and it applies also to the original design in Fig. 1—is that not all cumulative rates are attainable as the sum of post-crest (exponential) waves. Time t from the crest to the end of a wave implies a crest rate $R(t_{\text{crest}}) = \mathtt{P}^{-t/\mathtt{TSD}} \cdot \mathtt{BCR}$ and a number of lower channels (counting active wave channels and the base channel) $k = 1 + \lfloor t/\mathtt{TSD} \rfloor$. The cumulative rate at the crest time is then

$$\sum_{i=0}^{k} \mathtt{P}^{i} \cdot R(t_{\text{crest}}) \;=\; \frac{1 - \mathtt{P}^{\lfloor t/\mathtt{TSD} \rfloor + 2}}{1 - \mathtt{P}} \cdot \mathtt{P}^{-t/\mathtt{TSD}} \cdot \mathtt{BCR}.$$

Sweeping t over $[0, \infty)$ gives cumulative rates in the achievable set

$$\mathcal{A} \;=\; \bigcup_{i=1}^{\infty} \mathcal{A}_i, \qquad \text{where} \quad \mathcal{A}_i \;=\; \left[\mathtt{P} \cdot \frac{\mathtt{P}^{-(i+1)} - 1}{\mathtt{P}^{-1} - 1} \cdot \mathtt{BCR},\; \frac{\mathtt{P}^{-(i+1)} - 1}{\mathtt{P}^{-1} - 1} \cdot \mathtt{BCR} \right]. \tag{5}$$

This set is depicted in Fig. 3. When $\mathtt{SR}$ is in $\mathcal{A}$, t_{crest} is found by determining the value that results in cumulative rate $\mathtt{SR}$ at the crest; otherwise, when $\mathtt{SR}$ is in a gap in $\mathcal{A}$, $t_{\text{crest}} = \mathtt{TSD}$. This can be expressed as

$$t_{\text{crest}} \;=\; \mathtt{TSD} \cdot \max \left\{ \mathtt{N} - \log_{1/\mathtt{P}} \left(\frac{1 - \mathtt{P}}{1 - \mathtt{P}^{\mathtt{N}}} \cdot \frac{\mathtt{SR}}{\mathtt{BCR}} \right),\, 1 \right\}. \tag{6}$$

Note that $\mathtt{SR} \in \mathcal{A}_i$ implies $\mathtt{N} = i + 1$. Also, $\mathtt{SR}$ in the gap between $\mathcal{A}_i$ and $\mathcal{A}_{i+1}$ implies $\mathtt{N} = i + 1$.

The design expressed in (4) starts waves at rate 0 except for the rare values of $\mathtt{SR}$ that are not in $\mathcal{A}$. Therefore, it does not solve the difficulties of measuring MRTT when the top wave is joined, as described in Section 2. However, this design falls within a family of designs that includes the current best practice, so it is a good starting point.

3.2 Design 2: Start Waves at $\mathtt{BCR}$

To avoid problems in measuring MRTT, we would like the prototype wave to start at some strictly positive rate. Here we adjust Design 1 so that each wave channel starts at a minimum rate $\mathtt{BCR}$. (We will abandon this requirement when $\mathtt{SR}/\mathtt{BCR}$ is very small.) This is very similar to—though already somewhat better than—the design described in the current protocol specification [5]. The current best practice as described in Section 4 has further refinements.

Starting waves at $\mathtt{BCR}$ forces the crest rate to be lower and may reduce the necessary $\mathtt{N}$. We would like the crest time to be such that the rate at the crest plus the rates of $\mathtt{N}-2$ waves and the base channel below equals $\mathtt{SR}-\mathtt{BCR}$. For values of $\mathtt{SR}$ such that $\mathtt{SR}-\mathtt{BCR} \in \mathcal{A}_i$, there is such a crest time; otherwise, when $\mathtt{SR}-\mathtt{BCR}$ lies between $\mathcal{A}_i$ and $\mathcal{A}_{i+1}$ set $t_{\text{crest}} = \mathtt{TSD}$ and the starting rate will be greater than $\mathtt{BCR}$. This is achieved by choosing $\mathtt{N} = i+1$, or to avoid the intermediate step of placing $\mathtt{SR}-\mathtt{BCR}$ in or between the $\mathcal{A}_k$s,

$$\mathtt{N} \;=\; \left\lceil \log_{1/\mathtt{P}}\left(1 + \frac{1}{\mathtt{P}}\left(\frac{1}{\mathtt{P}} - 1\right)\left(\frac{\mathtt{SR}}{\mathtt{BCR}} - 1\right)\right)\right\rceil - 1. \tag{7}$$

Then,

$$t_{\text{crest}} \;=\; \mathtt{TSD} \cdot \max\left\{\mathtt{N} - \log_{1/\mathtt{P}}\left(\frac{1-\mathtt{P}}{1-\mathtt{P}^{\mathtt{N}}} \cdot \left(\frac{\mathtt{SR}}{\mathtt{BCR}} - 1\right)\right),\, 1\right\}. \tag{8}$$

Finally, the prototype wave is given by

$$R(t) \;=\; \begin{cases} \mathtt{BCR}, & 0 \le t < t_{\text{crest}} - \mathtt{TSD}; \\ \mathtt{SR} - \dfrac{\mathtt{P}^{-\mathtt{N}} - 1}{\mathtt{P}^{-1} - 1} \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR}, & t_{\text{crest}} - \mathtt{TSD} \le t < \mathtt{TSD}; \\ \mathtt{SR} - \left(1 + \dfrac{\mathtt{P}^{-(\mathtt{N}-1)} - 1}{\mathtt{P}(\mathtt{P}^{-1} - 1)} \cdot \mathtt{P}^{t/\mathtt{TSD}}\right) \cdot \mathtt{BCR}, & \mathtt{TSD} \le t < t_{\text{crest}}; \\ \mathtt{P}^{-\mathtt{N}+t/\mathtt{TSD}} \cdot \mathtt{BCR}, & t_{\text{crest}} \le t \le \mathtt{N} \cdot \mathtt{TSD}. \end{cases} \tag{9}$$

Design 2 makes sense for $\mathtt{SR}/\mathtt{BCR} \ge 2$; Design 1 can be used for smaller values of $\mathtt{SR}$. An example of the cumulative output rate is shown in Fig. 4.

The minimum rate in (9) occurs either at $t = 0^+$ or at $t = \mathtt{TSD}^+$. For the latter we have

$$R(\mathtt{TSD}^+) \;=\; \mathtt{SR} - \left(1 + \frac{\mathtt{P}^{-(\mathtt{N}-1)} - 1}{\mathtt{P}^{-1} - 1}\right) \cdot \mathtt{BCR} \;\ge\; \mathtt{P} \cdot \mathtt{BCR},$$

where the bound can be obtained by rearranging (7). Design 2 thus partially solves the problems associated with joining the top wave: Since the fluid rate is at least $\mathtt{P} \cdot \mathtt{BCR}$ when a join is issued, a flow of packets should begin immediately. This prevents join timeouts and spuriously very large MRTT measurements as large as $\mathtt{TSD}$.

A weakness remains in this design, however. Receiving at a high rate in WEBRC requires a small value of $\mathtt{ARTT} \cdot \sqrt{\mathtt{LOSSP}}$. Thus, presumably, the MRTT must be very small. When the top wave is joined and its rate is $\mathtt{BCR}$, the interpacket spacing is $1/\mathtt{BCR}$ seconds. If no modification is made to the MRTT calculation for the top wave, MRTT values will have deviation on the order of $1/\mathtt{BCR}$. This noise-like contribution to MRTT may be much larger than the true value, and large MRTTs will drive the target rate down. Lab testing and *ns* simulation of a design similar to Design 2 exhibited reception rates with large-amplitude limit cycles.[3]

[3] This testing used the wave design in [5], which differs slightly from Design 2 presented here.

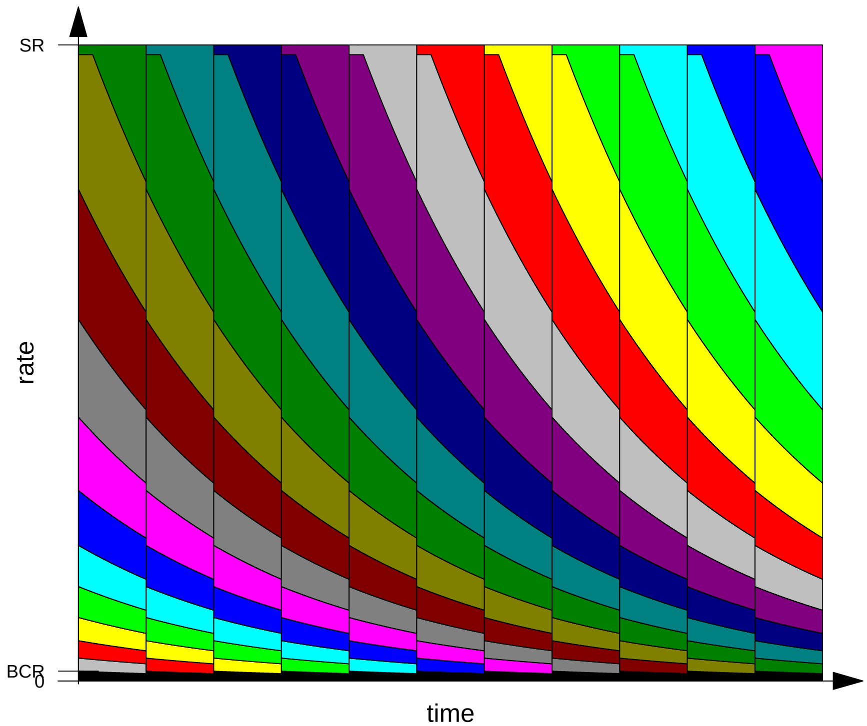

Figure 4: Example of Design 2 ($\mathtt{P} = \frac{3}{4}$, $\mathtt{SR}/\mathtt{BCR} = 65$).

3.3 Unification of Designs 1 and 2

It is worth noting that Designs 1 and 2 can be unified with a single set of equations. With $\mu = 0$ for Design 1 and $\mu = 1$ for Design 2, both of these designs are described as follows:

$$\mathtt{N} = \left\lceil \log_{1/\mathtt{P}} \left(1 + \frac{1}{\mathtt{P}}\left(\frac{1}{\mathtt{P}} - 1\right)\left(\frac{\mathtt{SR}}{\mathtt{BCR}} - \mu\right)\right)\right\rceil - 1 \tag{10}$$

$$t_{\text{crest}} = \mathtt{TSD} \cdot \max\left\{\mathtt{N} - \log_{1/\mathtt{P}}\left(\frac{1-\mathtt{P}}{1-\mathtt{P}^{\mathtt{N}}} \cdot \left(\frac{\mathtt{SR}}{\mathtt{BCR}} - \mu\right)\right), 1\right\} \tag{11}$$

$$R(t) = \begin{cases} \mu \cdot \mathtt{BCR}, & 0 \le t < t_{\text{crest}} - \mathtt{TSD}; \\ \mathtt{SR} - \dfrac{\mathtt{P}^{-\mathtt{N}} - 1}{\mathtt{P}^{-1} - 1} \cdot \mathtt{P}^{t/\mathtt{TSD}} \cdot \mathtt{BCR}, & t_{\text{crest}} - \mathtt{TSD} \le t < \mathtt{TSD}; \\ \mathtt{SR} - \left(\mu + \dfrac{\mathtt{P}^{-(\mathtt{N}-1)} - 1}{\mathtt{P}(\mathtt{P}^{-1} - 1)} \cdot \mathtt{P}^{t/\mathtt{TSD}}\right) \cdot \mathtt{BCR}, & \mathtt{TSD} \le t < t_{\text{crest}}; \\ \mathtt{P}^{-\mathtt{N}+t/\mathtt{TSD}} \cdot \mathtt{BCR}, & t_{\text{crest}} \le t \le \mathtt{N} \cdot \mathtt{TSD}. \end{cases} \tag{12}$$

In this unified design, the influence of μ is most transparent in $R(t)$ for $t \in [0, t_{\text{crest}} - \mathtt{TSD}]$: $\mu \cdot \mathtt{BCR}$ is the rate at the beginning of a wave, as long as $t_{\text{crest}} > \mathtt{TSD}$ is satisfied. When $t_{\text{crest}} = \mathtt{TSD}$, the rate at the beginning of a wave is larger, so $\mu \cdot \mathtt{BCR}$ is the minimum rate at the beginning of a wave. One modification pursued in Section 4 is to increase μ so that the interpacket spacing at the beginning of a wave is reduced.

The maximum interpacket spacing in the first two time slots of a wave depends on the minimum rate in that interval. The minimum rate may be at least $\mu \cdot \mathtt{BCR}$ or it may be $R(\mathtt{TSD}^+)$. Because of the way $\mathtt{N}$ and μ are related, one can prove

$$R(\mathtt{TSD}^{+}) \;=\; \mathtt{SR} - \left(\mu + \frac{\mathtt{P}^{-(\mathtt{N}-1)} - 1}{\mathtt{P}^{-1} - 1}\right) \cdot \mathtt{BCR} \;\geq\; \mathtt{P} \cdot \mathtt{BCR} \tag{13}$$

and that this bound is tight for some values of $\mathtt{SR}/\mathtt{BCR}$. The fact that the bound (13) is independent of μ suggests that some additional degree of freedom is needed to increase the minimum rate. In Section 4, a parameter is introduced to (10) that potentially changes $\mathtt{N}$ to decrease t_{crest} and increase $R(\mathtt{TSD}^{+})$.

While the minimum rate will be increased in Section 4, note that the minimum rate is generally less important than the starting rate because the minimum rate affects MRTT measurements for a narrower range of join times. With reference to Fig. 4, for example, note that the starting rate influences the variance of MRTT measurements made when the join time is in $[0, t_{\text{crest}} - \mathtt{TSD}]$ while the minimum rate affects MRTTs when the join time is in a small interval $[\mathtt{TSD}, \mathtt{TSD}+\epsilon]$. In addition to the sizes of these intervals, note that all receivers with target rates greater than $\mathtt{SR}$ have join times very close to 0.

4 Current Wave Design

The weakness of Design 2 described above suggests that performance can be improved simply by increasing μ. While it is true that this will improve (lower) the interpacket spacing at the beginning of a wave, excessively large values of μ are not desirable. Larger values of μ shorten the exponential portion of the wave, and it is the exponential portion that allows fine-grained rate control. In the extreme, making μ close to $\mathtt{SR}/\mathtt{BCR}$ puts almost all of the packets into the (non-exponential) first time slots of the waves.

The exponential shape of the waves, and hence the manner of rate increase as waves are joined, suggests a good way to choose μ. The fraction of the total rate on the top wave is approximately $1 - \Gamma^{-1} \approx 1 - \mathtt{P}$. This fraction being approximately fixed suggests that μ should grow in proportion to $\mathtt{SR}/\mathtt{BCR}$ and certainly not be larger than $(1 - \mathtt{P})(\mathtt{SR}/\mathtt{BCR})$. Specifically, we use

$$\mu \;=\; \psi \cdot (1 - \mathtt{P}) \left(\frac{\mathtt{SR}}{\mathtt{BCR}} - 1\right), \tag{14}$$

where $\psi \in (0, 1)$ trades off the interpacket spacing at the beginning of the wave against the length of the exponential portion of the wave; its value is chosen later. Using an affine function of $\mathtt{SR}/\mathtt{BCR}$ rather than a linear function makes it possible to apply a single set of equations for all values of $\mathtt{SR}/\mathtt{BCR} \geq 1$. (Recall that Design 2 does not make sense for $\mathtt{SR}/\mathtt{BCR} < 2$.)

Increasing μ while leaving all other computations in (10)–(12) unchanged does not always increase the minimum rate of a wave, as shown in (13). We now alter the design to increase the minimum rate. First, it helps to understand how $\mathtt{N}$, t_{crest} and $R(\mathtt{TSD}^{+})$ vary when they are determined through (10)–(12). Let μ be given by (14) with $\psi = \frac{1}{2}$. Then μ, $\mathtt{N}$, t_{crest} and $R(\mathtt{TSD}^{+})$ vary with $\mathtt{SR}/\mathtt{BCR}$ as shown in Fig. 5. This plot confirms that the bound in (13) is tight. The smallest values of $R(\mathtt{TSD}^{+})$ occur for the $\mathtt{SR}$ values just above a threshold where $\mathtt{N}$ has changed. At these critical values of $\mathtt{SR}$, $\mathtt{N}$ given by (10) can be replaced by its upper envelope

$$\mathtt{N}_{\text{upper}} \;=\; \log_{1/\mathtt{P}}\left(1 + \frac{1}{\mathtt{P}}\left(\frac{1}{\mathtt{P}} - 1\right)\left(\frac{\mathtt{SR}}{\mathtt{BCR}} - \mu\right)\right).$$

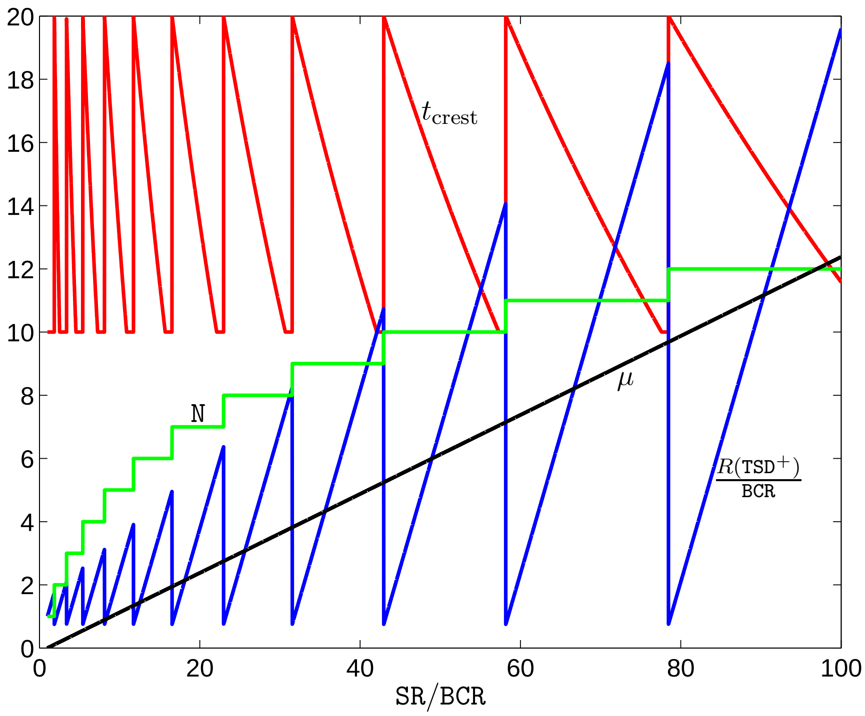

Figure 5: Variation of parameters through (10)–(12) when μ is given by (14). In this example, $\psi = \frac{1}{2}$, $\texttt{P} = \frac{3}{4}$ and $\texttt{TSD} = 10$.

Substituting

$$\texttt{P}^{\texttt{N}_{\text{upper}}} \;=\; \left(1 + \frac{1}{\texttt{P}}\left(\frac{1}{\texttt{P}} - 1\right)\left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right)\right)^{-1}$$

for $\texttt{P}^{\texttt{N}}$ in (11) and simplifying shows that the upper envelope of t_{crest} is $2 \cdot \texttt{TSD}$. Similarly, substituting in (12) gives the lower envelope of $R(\texttt{TSD}^+)$ as $\texttt{P} \cdot \texttt{BCR}$. These calculations explain some of the envelopes observed in Fig. 5 and the bound (13).

To increase the minimum rate by increasing $R(\texttt{TSD}^+)$, introduce a parameter $\varphi \in [0, 1]$ to (10):

$$\texttt{N} \;=\; \left\lceil \log_{1/\texttt{P}}\left(1 + \frac{1}{\texttt{P}^{1-\varphi}}\left(\frac{1}{\texttt{P}} - 1\right)\left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right)\right)\right\rceil - 1. \tag{15}$$

The influence of φ on t_{crest} is most transparent. Computing with the upper envelope of the new value of $\texttt{N}$ gives

$$\frac{1 - \texttt{P}}{1 - \texttt{P}^{\texttt{N}}} \cdot \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right) \;=\; \frac{1 - \texttt{P}}{1 - \texttt{P}^{\texttt{N}}} \cdot \frac{\texttt{P}^{1-\varphi}(\texttt{P}^{-\texttt{N}} - 1)}{\texttt{P}^{-1} - 1} \;=\; \texttt{P}^{-\texttt{N}+2-\varphi}.$$

Thus,

$$\texttt{N} - \log_{1/\texttt{P}}\left(\frac{1 - \texttt{P}}{1 - \texttt{P}^{\texttt{N}}} \cdot \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right)\right) \;=\; \texttt{N} - (\texttt{N} - 2 + \varphi) \;=\; 2 - \varphi$$

and so the upper envelope of t_{crest} becomes $(2 - \varphi) \cdot \texttt{TSD}$.

Though t_{crest} is visually striking, we are more interested in the effect of φ on the minimum rate. Using the upper envelope of (15), we have the following:

$$\begin{aligned}
\frac{R(\texttt{TSD}^+)}{\texttt{BCR}} &= \frac{\texttt{SR}}{\texttt{BCR}} - \mu - \frac{\texttt{P}^{-(\texttt{N}-1)} - 1}{\texttt{P}(\texttt{P}^{-1} - 1)} \cdot \texttt{P} \\
&= \frac{\texttt{SR}}{\texttt{BCR}} - \mu - \frac{1}{\texttt{P}^{-1} - 1} \left[\texttt{P} \cdot \texttt{P}^{-\texttt{N}} - 1\right] \\
&= \frac{\texttt{SR}}{\texttt{BCR}} - \mu - \frac{1}{\texttt{P}^{-1} - 1} \left[\texttt{P} \cdot \left(1 + \frac{1}{\texttt{P}^{1-\varphi}} \left(\frac{1}{\texttt{P}} - 1\right) \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right)\right) - 1\right] \\
&= \frac{\texttt{SR}}{\texttt{BCR}} - \mu - \frac{1}{\texttt{P}^{-1} - 1} \left[\frac{1}{\texttt{P}^{1-\varphi}}(1 - \texttt{P}) \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right) - (1 - \texttt{P})\right] \\
&= (1 - \texttt{P}^{\varphi}) \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right) + \texttt{P} \\
&= (1 - \texttt{P}^{\varphi}) \left(\frac{\texttt{SR}}{\texttt{BCR}} - \psi(1 - \texttt{P}) \left(\frac{\texttt{SR}}{\texttt{BCR}} - 1\right)\right) + \texttt{P} \\
&= (1 - \texttt{P}^{\varphi}) \left(1 - \psi(1 - \texttt{P})\right) \left(\frac{\texttt{SR}}{\texttt{BCR}} - 1\right) + 1 - \texttt{P}^{\varphi} + \texttt{P} \qquad (16)
\end{aligned}$$

Thus introducing $\varphi > 0$ makes $R(\texttt{TSD}^+)$ approximately a linear function of $\texttt{SR}$. Considering also $R(0^+)$, the minimum rate can be bounded below by an affine function of $\texttt{SR}$ with slope

$$\min\{\psi(1 - \texttt{P}),\ (1 - \texttt{P}^{\varphi})\,(1 - \psi(1 - \texttt{P}))\}.$$

As an example, Fig. 6 shows the same variables as Fig. 5 but with $\varphi = \frac{1}{3}$. This confirms that the upper envelope of t_{crest} is $(2-\varphi) \cdot \texttt{TSD}$ and that the lower envelope of $R(\texttt{TSD}^+)$ is the affine function given by (16).

We are now prepared to finalize the wave design by choosing ψ and φ. As mentioned previously, increasing ψ shortens the exponential portion of the wave. The effect of increasing φ is not quite as obvious. Increasing φ decreases t_{crest}—thus seemingly lengthening the exponential portion of the wave—but also may decrease $\texttt{N}$.

When a wave crests, its rate is $\texttt{P}^{-\texttt{N}+t_{\text{crest}}/\texttt{TSD}} \cdot \texttt{BCR}$. The cumulative rate of the crested wave and all lower channels is

$$R_{\text{crest}} = \sum_{i=0}^{N-1} \texttt{P}^{-i} \cdot \texttt{P}^{-\texttt{N}+t_{\text{crest}}/\texttt{TSD}} \cdot \texttt{BCR} = \frac{1 - \texttt{P}^{\texttt{N}}}{1 - \texttt{P}} \cdot \texttt{P}^{-\texttt{N}+t_{\text{crest}}/\texttt{TSD}} \cdot \texttt{BCR} = \frac{\texttt{P}^{-\texttt{N}} - 1}{1 - \texttt{P}} \cdot \texttt{P}^{t_{\text{crest}}/\texttt{TSD}} \cdot \texttt{BCR}.$$

Fine-grained rate control is possible up to rate R_{crest}. Fig. 7 shows R_{crest} as a fraction of $\texttt{SR}$, as a function of $\texttt{SR}$, for a few (ψ, φ) pairs. It appears that upper and lower envelopes both converge to constants. This is in fact true, and we now determine these constants.

Using the upper envelope of $\texttt{N}$ and upper envelope of t_{crest} gives

$$\begin{aligned}
R_{\text{crest}} &= \frac{\left(1 + \frac{1}{\texttt{P}^{1-\varphi}} \left(\frac{1}{\texttt{P}} - 1\right) \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right)\right) - 1}{1 - \texttt{P}} \cdot \texttt{P}^{2-\varphi} \cdot \texttt{BCR} = \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right) \cdot \texttt{BCR} = \texttt{SR} - \mu \cdot \texttt{BCR} \\
&= \texttt{SR} - \psi(1 - \texttt{P}) \left(\frac{\texttt{SR}}{\texttt{BCR}} - 1\right) \cdot \texttt{BCR} = (1 - \psi(1 - \texttt{P})) \cdot \texttt{SR} + \psi(1 - \texttt{P}) \cdot \texttt{BCR}.
\end{aligned}$$

This shows that the upper envelope of $R_{\text{crest}}/\texttt{SR}$ in Fig. 7 approaches the constant value $1-\psi(1-\texttt{P})$. The lower envelope of $\texttt{N}$ and lower envelope of t_{crest} give

$$R_{\text{crest}} = \frac{\texttt{P} \cdot \left(1 + \frac{1}{\texttt{P}^{1-\varphi}} \left(\frac{1}{\texttt{P}} - 1\right) \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right)\right) - 1}{1 - \texttt{P}} \cdot \texttt{P} \cdot \texttt{BCR} = \left(\frac{1}{\texttt{P}^{1-\varphi}} \left(\frac{\texttt{SR}}{\texttt{BCR}} - \mu\right) - 1\right) \cdot \texttt{P} \cdot \texttt{BCR}$$

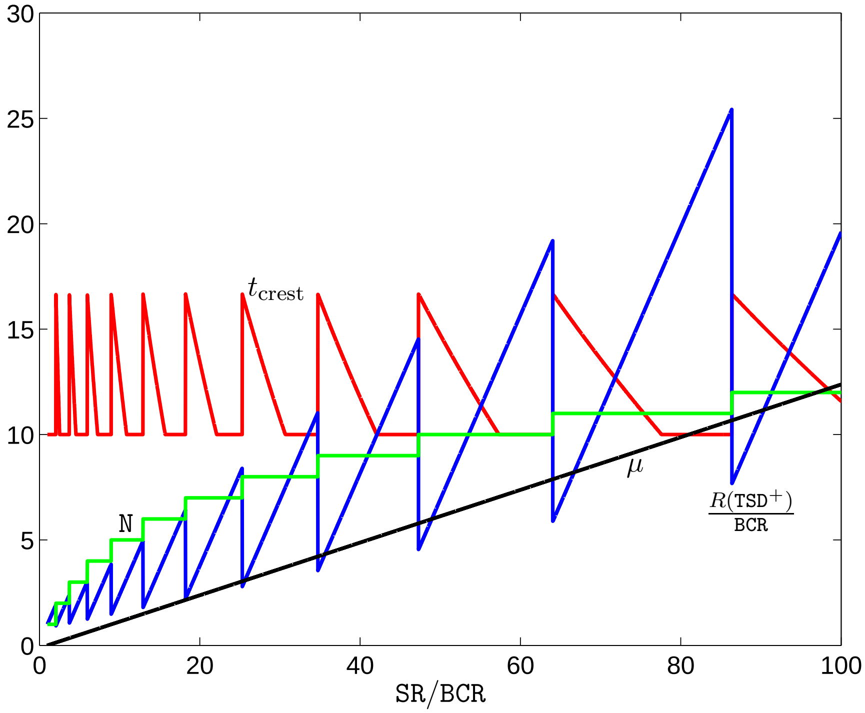

Figure 6: Variation of parameters through (15), (11) and (12) when μ is given by (14). In this example, $\psi = \frac{1}{2}$, $\phi = \frac{1}{3}$, $\texttt{P} = \frac{3}{4}$ and $\texttt{TSD} = 10$.

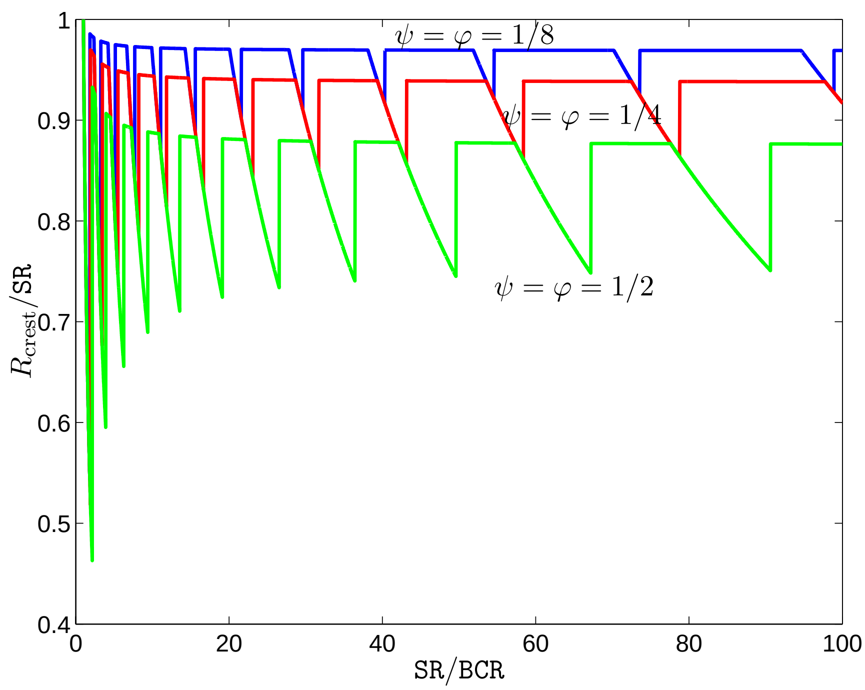

Figure 7: Fraction of $\texttt{SR}$ for which fine-grained rate control is possible ($\texttt{P} = \frac{3}{4}$).

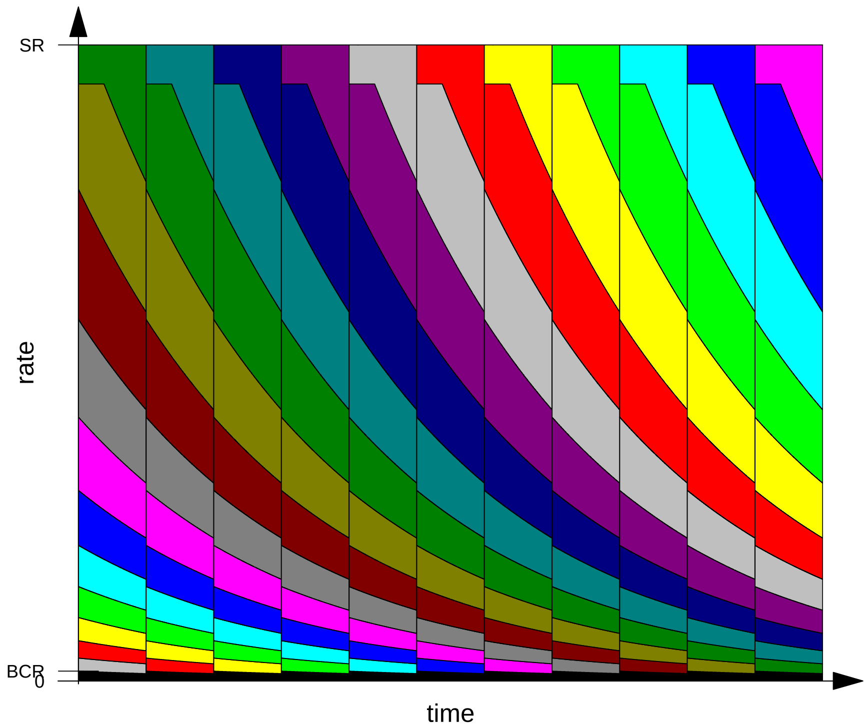

Figure 8: Example of the final design ($\mathtt{P} = \frac{3}{4}$, $\psi = \frac{1}{4}$, $\phi = \frac{1}{5}$, $\mathtt{SR}/\mathtt{BCR} = 65$).

$$= \quad \mathtt{P}^{\varphi}(\mathtt{SR} - \mu \cdot \mathtt{BCR}) - \mathtt{P} \cdot \mathtt{BCR} \;=\; \mathtt{P}^{\varphi}(1 - \psi(1 - \mathtt{P})) \cdot \mathtt{SR} + (\mathtt{P}^{\varphi}\psi(1 - \mathtt{P}) - \mathtt{P}) \cdot \mathtt{BCR}.$$

Thus the lower envelope of $R_{\text{crest}}/\mathtt{SR}$ approaches $\mathtt{P}^{\varphi}(1 - \psi(1 - \mathtt{P}))$. The upper and lower envelope calculations are consistent with Fig. 7.

To summarize, to have low interpacket spacing at the beginning of each wave, we would like for

$$\psi(1 - \mathtt{P}) \qquad \text{and} \qquad (1 - \mathtt{P}^{\varphi})(1 - \psi(1 - \mathtt{P}))$$

to be large. To have fine-grained rate control for a large range of rates, we would like for

$$1 - \psi(1 - \mathtt{P}) \qquad \text{and} \qquad \mathtt{P}^{\varphi}(1 - \psi(1 - \mathtt{P}))$$

to be large. These objectives thus stand in direct opposition and any choice of ψ and ϕ can be second-guessed. The design space can be limited somewhat by the following observation: There is no point in making $(1 - \mathtt{P}^{\varphi})(1 - \psi(1 - \mathtt{P}))$ larger than $\psi(1 - \mathtt{P})$ because the rate at the beginning of a wave is more important than the minimum over the first two time slots. This translates to limiting φ to

$$\varphi \leq \log_{1/\mathtt{P}} \left(\frac{1 - \psi(1 - \mathtt{P})}{1 - 2\psi(1 - \mathtt{P})} \right). \tag{17}$$

On the other hand, φ should not be much smaller than the value in (17) because pathologies in MRTT will then be dominated by small $R(\mathtt{TSD}^{+})$ values. Furthermore, with reference to Fig. 7 note that the negative impact of large φ values on granularity of rate control is somewhat overstated by looking only at envelopes.

The suggested values for ψ and ϕ to use with the default value of $\mathtt{P} = \frac{3}{4}$ are $\frac{1}{4}$ and $\frac{1}{5}$, respectively. An example of the resulting design is given in Fig. 8.

5 Computing Transmission Times from the Fluid Model

This section describes a method for converting the fluid model rate given by (12) to a set of packet transmission times. According to the fluid model, the number of packets sent in one time slot is $K = \mathtt{SR} \cdot \mathtt{TSD}$. We henceforth assume that this quantity is integral. This allows the scheduling of packet transmissions on all channels in all time slots to be described in terms of a computation involving one period of the base channel and one wave.

Form a function $f(t)$, $t \in [0, (\mathtt{N}+1) \cdot \mathtt{TSD}]$, by taking one period of the base channel followed by one time-reversed wave. That is,

$$f(t) = \begin{cases} R_{\text{base}}(t), & t \in [0,\ \mathtt{TSD}]; \\ R((\mathtt{N}+1) \cdot \mathtt{TSD} - t), & t \in (\mathtt{TSD},\ (\mathtt{N}+1) \cdot \mathtt{TSD}]. \end{cases}$$

This function is strictly positive on $[0, (\mathtt{N}+1) \cdot \mathtt{TSD}]$ and integrates over this interval to $\mathtt{SR} \cdot \mathtt{TSD}$. Choosing times at which to transmit packets is then equivalent to dividing the area under $f(t)$ into $\mathtt{SR} \cdot \mathtt{TSD}$ regions of unit area. The left edge of each of these regions is taken as a packet transmission time. An example of this graphical construction is given in Fig. 9. The algebraic equivalent is to solve

$$\int_0^{\tau_k} f(t)\, dt = k, \qquad k = 0,\ 1,\ 2,\ \ldots,\ \mathtt{SR} \cdot \mathtt{TSD} - 1. \tag{18}$$

Though it may seem arbitrary, there are reasons for taking the base channel period first, using the time-reversed wave, and equating left edges of regions to transmission times. Taking the base channel first makes the first packet of each time slot a base channel packet; thus, every receiver, regardless of rate, learns of time slot boundaries as soon as possible. Using the time-reversed wave causes the first packet of a wave to occur as late as possible while being consistent with the fluid model; this is important because receivers cannot join a wave until the indication of a new time slot reveals that the wave has started. Finally, with this configuration it is possible to precompute the τ_ks that depend on TSD, BCR, and P but are independent of SR.

The first solution of (18) is $\tau_0 = 0$, so the first base channel packet is sent at time $b_0 = 0$. The next is sent at the time b_1 that satisfies

$$1 \;=\; \int_0^{b_1} R_{\text{base}}(\tau)\, d\tau \;=\; \int_0^{b_1} \mathtt{P}^{\tau/\mathtt{TSD}} \cdot \mathtt{BCR}\, d\tau.$$

Continuing this process, a total of $L = \lceil \frac{\mathtt{BCR} \cdot \mathtt{TSD}(1-\mathtt{P})}{-\ln \mathtt{P}} \rceil$ packets are sent at times

$$b_k \;=\; \mathtt{TSD} \cdot \log_{\mathtt{P}} \left(1 + \frac{\ln \mathtt{P}}{\mathtt{BCR} \cdot \mathtt{TSD}} \cdot k \right), \qquad k = 0,\ 1,\ \ldots, L-1. \tag{19}$$

Conceptually, this process is continued to find all of the wave channel transmission times. The key differences are integrating from the end of the wave rather than the beginning and accounting for the difference between the number of base channel packets per time slot according to the fluid model,

$$I_0 \;=\; \int_0^{\mathtt{TSD}} R_{\text{base}}(\tau)\, d\tau \;=\; \int_0^{\mathtt{TSD}} \mathtt{P}^{\tau/\mathtt{TSD}} \cdot \mathtt{BCR}\, d\tau \;=\; \frac{\mathtt{BCR} \cdot \mathtt{TSD}(1-\mathtt{P})}{-\ln \mathtt{P}},$$

and the number of transmitted packets $L = \lceil I_0 \rceil$.

Let

$$S(s) \;=\; \int_{\mathtt{N} \cdot \mathtt{TSD} - s}^{\mathtt{N} \cdot \mathtt{TSD}} R(\tau)\, d\tau.$$

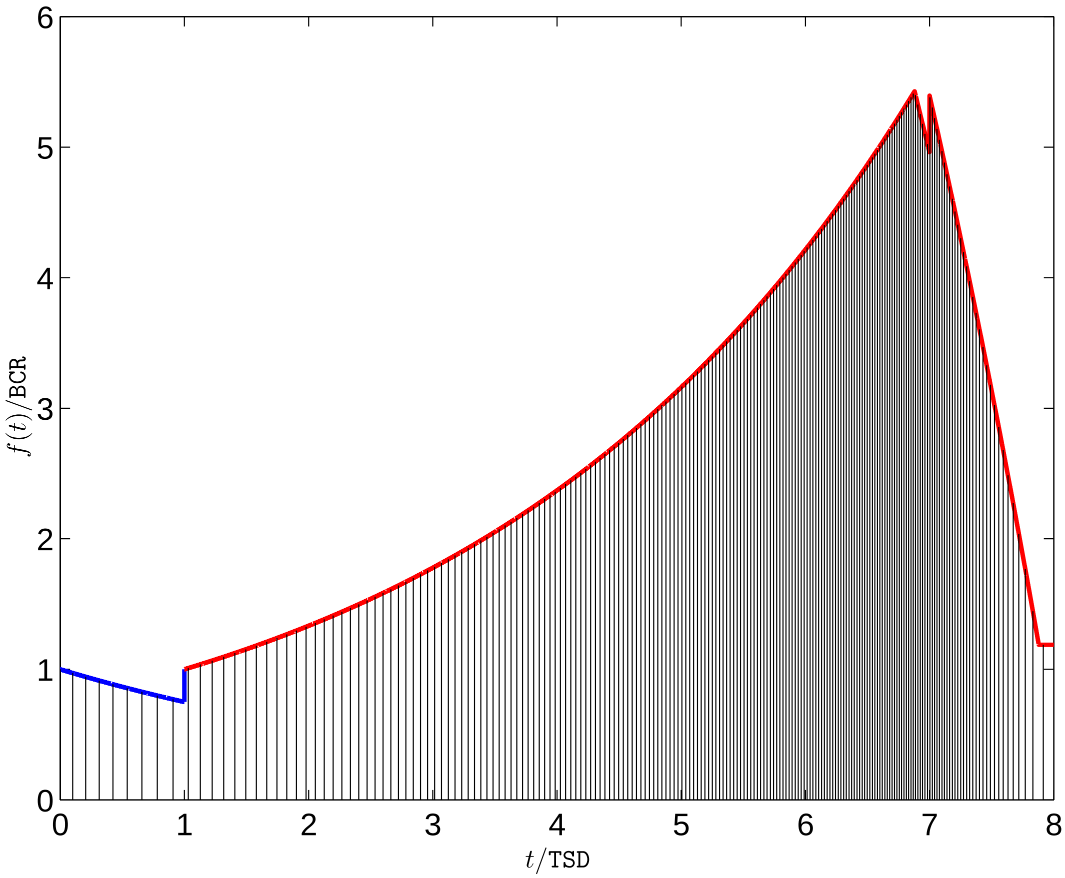

Figure 9: Graphical depiction of the computation of packet transmission times from fluid-model rates. The packet transmission times divide the area under $f(t)$ into $\mathtt{SR} \cdot \mathtt{TSD}$ unit-area regions. In this example $\mathtt{P} = \frac{3}{4}$, $\mathtt{BCR} = 1$, $\mathtt{SR} = 20$, and $\mathtt{TSD} = 10$.

(s is time running backward from the end of the wave.) $S(s)$ is the fluid-model approximation to the number of packets to be sent in the last s seconds of the wave. The desired transmission times (measured backward from the end of the wave) are the solutions to

$$I_0 + S(s_k) \;=\; k, \qquad k = L,\, L+1,\, \ldots,\, K-1. \tag{20}$$

Solving (20) for the s_ks is complicated by the fact that $R(t)$, $t \in [0, \mathtt{N} \cdot \mathtt{TSD}]$ is described separately on four subintervals (see (12)). The remainder of this section details the computations of the s_ks.

For the last time range in (12), the integral of $R(t)$ has a simple form:

$$S(s) \;=\; \frac{\mathtt{BCR} \cdot \mathtt{TSD}}{-\ln \mathtt{P}} \left(\mathtt{P}^{-s/\mathtt{TSD}} - 1\right) \qquad \text{for } s \in [0, \mathtt{N} \cdot \mathtt{TSD} - t_{\text{crest}}].$$

Thus for $k = L$ and some larger values of k, (20) becomes

$$I_0 + \frac{\mathtt{BCR} \cdot \mathtt{TSD}}{-\ln \mathtt{P}} \left(\mathtt{P}^{-s_k/\mathtt{TSD}} - 1\right) \;=\; k. \tag{21}$$

The range of k for which (21) is valid is determined based on the number of fluid-model base channel packets I_0 and the number of fluid-model packets sent in the last $\mathtt{N} \cdot \mathtt{TSD} - t_{\text{crest}}$ seconds of a wave. The latter quantity is given by

$$I_1 \;=\; \int_{t_{\text{crest}}}^{\texttt{N}\cdot\texttt{TSD}} R(\tau)\, d\tau \;=\; S(\texttt{N}\cdot\texttt{TSD} - t_{\text{crest}}) \;=\; \frac{\texttt{BCR}\cdot\texttt{TSD}}{-\ln \texttt{P}} \left(\texttt{P}^{-\texttt{N}+t_{\text{crest}}/\texttt{TSD}} - 1\right).$$

The sum of the number base channel packets in one time slot and the number of wave channel packets in the portion of the wave under consideration is denoted $K_1 = \lceil I_0 + I_1 \rceil$. Thus (21) is solved to obtain

$$s_k \;=\; \texttt{TSD}\cdot \log_{1/\texttt{P}}\left(1 - \frac{\ln \texttt{P}}{\texttt{BCR}\cdot\texttt{TSD}}\,(k - I_0)\right), \qquad k = L,\, L+1,\, \ldots,\, K_1 - 1. \tag{22}$$

Computing and inverting $S(s)$ gets uglier from here on. For $s \in (\texttt{N}\cdot\texttt{TSD} - t_{\text{crest}},\, (\texttt{N}-1)\cdot\texttt{TSD}]$, $S(s)$ has a fixed component I_1 plus

$$\begin{aligned}\int_{t_{\text{crest}}-\tau}^{t_{\text{crest}}} R(t)\, dt \;&=\; \int_{t_{\text{crest}}-\tau}^{t_{\text{crest}}} \left(\texttt{SR} - \mu\cdot\texttt{BCR} - \frac{\texttt{P}^{-(\texttt{N}-1)} - 1}{\texttt{P}(\texttt{P}^{-1} - 1)}\cdot \texttt{P}^{t/\texttt{TSD}}\cdot\texttt{BCR}\right) dt \\ &=\; \underbrace{(\texttt{SR} - \mu\cdot\texttt{BCR})}_{A}\,\tau - \underbrace{\frac{\texttt{P}^{-(\texttt{N}-1)} - 1}{\texttt{P}(\texttt{P}^{-1} - 1)}\cdot\texttt{BCR}\cdot\frac{\texttt{TSD}}{-\ln \texttt{P}}\cdot \texttt{P}^{t_{\text{crest}}/\texttt{TSD}}}_{B}\left(\texttt{P}^{-\tau/\texttt{TSD}} - 1\right),\end{aligned}$$

where $\tau = s - (\texttt{N}\cdot\texttt{TSD} - t_{\text{crest}})$. The total fluid contribution of the latter for this portion of the wave is

$$I_2 \;=\; \int_{\texttt{TSD}}^{t_{\text{crest}}} R(t)\, dt \;=\; A(t_{\text{crest}} - \texttt{TSD}) - B\left(\texttt{P}^{-(t_{\text{crest}}-\texttt{TSD})/\texttt{TSD}} - 1\right).$$

Thus the number of packets for one period of the base channel plus the number of packets for the last $(\texttt{N}-1)\texttt{TSD}$ seconds of a wave is $K_2 = \lceil I_0 + I_1 + I_2 \rceil$ and scheduling the packets for $s \in (\texttt{N}\cdot\texttt{TSD} - t_{\text{crest}}, (\texttt{N}-1)\texttt{TSD}]$ amounts to solving

$$I_0 + I_1 + A\tau_k - B(\texttt{P}^{-\tau_k/\texttt{TSD}} - 1) \;=\; k, \qquad k = K_1,\, K_1 + 1, \ldots, K_2 - 1 \tag{23}$$

for the τ_ks. There is no elementary closed form solution; techniques for finding approximate solutions are described in Appendix A. The τ_ks in this calculation are transmission times measured backward in time from time t_{crest}. We thus obtain $s_k = \tau_k + \texttt{N}\cdot\texttt{TSD} - t_{\text{crest}}$, $k = K_1,\, K_1+1,\, \ldots,\, K_2-1$.

The calculations are very similar for the next interval, $s \in ((\texttt{N}-1)\texttt{TSD}, (\texttt{N}+1)\texttt{TSD} - t_{\text{crest}}]$. We are interested in solutions to

$$S(\tau + (\texttt{N}-1)\texttt{TSD}) \;=\; I_0 + I_1 + I_2 + \int_{\texttt{TSD}-\tau}^{\texttt{TSD}} R(t)\, dt \;=\; k \qquad \text{for } \tau \in [0,\, 2\cdot\texttt{TSD} - t_{\text{crest}}].$$

Integrating yields

$$I_0 + I_1 + I_2 + A'\tau_k - B'(\texttt{P}^{-\tau_k/\texttt{TSD}} - 1) \;=\; k \tag{24}$$

where $A' = \texttt{SR}$ and

$$B' \;=\; \frac{\texttt{BCR}\cdot\texttt{TSD}}{-\ln \texttt{P}}\cdot\frac{\texttt{P}^{-(\texttt{N}-1)} - \texttt{P}}{\texttt{P}^{-1} - 1}.$$

The range of integers k for which we solve (24) depends on how many packets are transmitted in this time interval. Paralleling the previous calculations, let

$$I_3 \;=\; \int_{t_{\text{crest}}-\texttt{TSD}}^{\texttt{TSD}} R(t)\, dt \;=\; \texttt{SR}(2\cdot\texttt{TSD} - t_{\text{crest}}) - \frac{\texttt{BCR}\cdot\texttt{TSD}}{-\ln \texttt{P}}\cdot\frac{\texttt{P}^{-(\texttt{N}-1)} - \texttt{P}}{\texttt{P}^{-1} - 1}\cdot\left(\texttt{P}^{t_{\text{crest}}/\texttt{TSD}-2} - 1\right)$$

and $K_3 = \lceil I_0 + I_1 + I_2 + I_3 \rceil$. Then (24) is solved for $k = K_2, K_2 + 1, \ldots, K_3 - 1$. Of course, (24) is of the same form as (23) and solutions can be approximated similarly. This process yields $s_k = \tau_k + (\mathtt{N} - 1) \cdot \mathtt{TSD}$, $k = K_2, K_2 + 1, \ldots, K_3 - 1$.

In the final interval $s \in ((\mathtt{N}+1)\mathtt{TSD} - t_{\text{crest}}, \mathtt{N} \cdot \mathtt{TSD}]$, the fluid-model rate is $\mu \cdot \mathtt{BCR}$ so the time between packets is $(\mu \cdot \mathtt{BCR})^{-1}$. To correctly patch together the intervals, we seek solutions to

$$I_0 + I_1 + I_2 + I_3 + \mu \cdot \mathtt{BCR} \cdot \tau_k \;=\; k, \qquad k = K_3, K_3 + 1, \ldots, K - 1. \tag{25}$$

The solutions are

$$\tau_k \;=\; \frac{1}{\mu \cdot \mathtt{BCR}} \left(k - I_0 - I_1 - I_2 - I_3\right)$$

and yield

$$s_k \;=\; \frac{1}{\mu \cdot \mathtt{BCR}} \left(k - I_0 - I_1 - I_2 - I_3\right) + (\mathtt{N}+1) \cdot \mathtt{TSD} - t_{\text{crest}}, \qquad k = K_3, K_3 + 1, \ldots, K - 1. \tag{26}$$

As a sanity check for these calculations, let

$$I_4 \;=\; \int_0^{t_{\text{crest}} - \mathtt{TSD}} R(t)\, dt \;=\; \mu \cdot \mathtt{BCR} \cdot (t_{\text{crest}} - \mathtt{TSD})$$

and verify $I_0 + I_1 + I_2 + I_3 + I_4 = \mathtt{TSD} \cdot \mathtt{SR}$.

6 Suggested Sender Implementation

Evaluating (19), (22), and (26) and solving (23) and (24) gives values over the real numbers as the times at which to transmit packets. Since we intend to have constant rate output and packet transmissions could turn out to be bursty for reasons unrelated to WEBRC, these times should be used only to determine an ordering of packets for one period of TSD seconds and then discarded. This ordering is determined by representing each transmission time as a quotient and remainder modulo TSD and then sorting the remainders. This is detailed below.

Recall that in each time slot the sender sends L packets on the base channel at times $\{b_k\}_{k=0}^{L-1}$ given by (19). Represent these times as (channel,time) pairs $\{(\mathtt{T}, b_k)\}_{k=0}^{L-1}$.

Without loss of generality, suppose the wave channel for which we have made computations is channel $\mathtt{N}-1$. This means that time 0 is the beginning of a time slot numbered 0. The wave channel packet transmission time s_k (measured from the end of the wave) is converted to a (channel,time) pair as follows. Let $\mathtt{N} \cdot \mathtt{TSD} - s_k = m \cdot \mathtt{TSD} + w_k$ for $m \in \mathbb{Z}$ and $w_k \in [0, \mathtt{TSD})$. Transmission of a packet on wave channel $\mathtt{N} - 1$ at time $\mathtt{N} \cdot \mathtt{TSD} - s_k$ implies a packet is sent on wave channel $\mathtt{N} - 1 - m \bmod \mathtt{T}$ at time w_k. Thus we have the mapping of s_k to $(\mathtt{N} - 1 - m \bmod \mathtt{T}, w_k)$.

The b_ks and w_ks are all in the interval $[0, \mathtt{TSD})$. By sorting the (channel,time) pairs by time, we get a sequence of transmission events for time slot 0. The same sequence can be used in time slot $i \neq 0$ if all channels $\mathtt{CN} \neq \mathtt{T}$ are replaced by $\mathtt{CN} + i \bmod \mathtt{T}$. There is no need to retain the times in the (channel,time) pairs after the sorting.

The transmissions on all the channels for a set of consecutive time slots is depicted in Fig. 10.

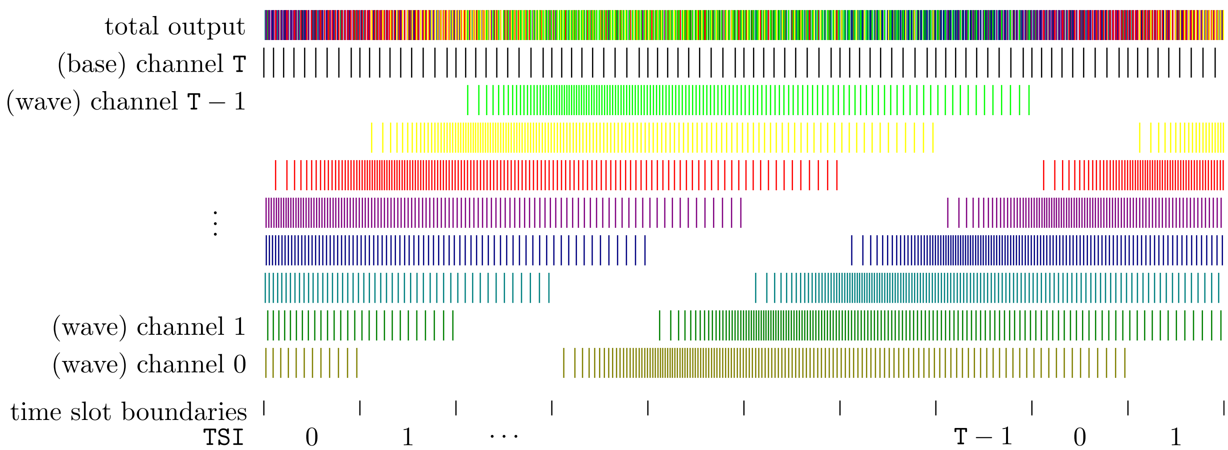

Figure 10: Idealized packet transmission times for a session with $\mathtt{BCR} = 1$, $\mathtt{SR} = 14$, $\mathtt{P} = 0.75$, $\mathtt{Q} = 2$, and $\mathtt{TSD} = 10$. (Larger ratios $\mathtt{SR}/\mathtt{BCR}$ are of more interest, and in practice $\mathtt{Q}$ is much larger.) Notice that waves start at time slot boundaries and reach their crests in their second active time slot. Taking all the channels together, the sender transmits at a constant rate.

A Solving $A\tau - B(\mathtt{P}^{-\tau/\mathtt{TSD}} - 1) = C$

Although there is no elementary closed-form solution, any number of root finding methods could be used to solve

$$A\tau - B(\mathtt{P}^{-\tau/\mathtt{TSD}} - 1) \;=\; C. \tag{27}$$

To have a formula (rather than a recursion) for τ, one can replace (27) with an approximate equation that has an elementary closed-form solution. For this purpose it is important to note that we are interested in $\tau \in [0, \mathtt{TSD})$. Thus we get reasonable accuracy with Taylor approximations of $\mathtt{P}^{-\tau/\mathtt{TSD}}$ about $\tau = 0$.

A first-order Taylor approximation is

$$\mathtt{P}^{-\tau/\mathtt{TSD}} \;=\; e^{-\tau(\ln \mathtt{P})/\mathtt{TSD}} \approx 1 - \frac{\ln \mathtt{P}}{\mathtt{TSD}}\tau.$$

This turns (27) into a linear equation with approximate solution

$$\tau \approx \frac{C}{A + B\frac{\ln \mathtt{P}}{\mathtt{TSD}}}. \tag{28}$$

A second-order Taylor approximation is

$$\mathtt{P}^{-\tau/\mathtt{TSD}} \;=\; e^{-\tau(\ln \mathtt{P})/\mathtt{TSD}} \approx 1 - \frac{\ln \mathtt{P}}{\mathtt{TSD}}\tau + \frac{1}{2}\left(\frac{\ln \mathtt{P}}{\mathtt{TSD}}\right)^2 \tau^2,$$

which yields

$$\underbrace{-\frac{B}{2}\left(\frac{\ln \mathtt{P}}{\mathtt{TSD}}\right)^2}_{a} \tau^2 + \underbrace{\left(A + B\frac{\ln \mathtt{P}}{\mathtt{TSD}}\right)}_{b} \tau - \underbrace{C}_{-c} \;=\; 0.$$

This is a standard quadratic equation; the desired solution is

$$\tau = \frac{-b+\sqrt{b^2-4ac}}{2a}.$$

The process of increasing the degree of the Taylor expansion can be continued to degree four; beyond the quadratic case there is no closed-form solution to the approximate polynomial equation that is obtained.

Since $\tau/\mathtt{TSD}$ is not always small enough for fast convergence of the Taylor series about 0, a better way to improve accuracy is to use Taylor approximations about different points. Let τ_0 be given by (28) and expand $f(\tau) = \mathtt{P}^{-\tau/\mathtt{TSD}}$ about τ_0 to get

$$\begin{aligned}
\mathtt{P}^{-\tau/\mathtt{TSD}} &= e^{-\tau(\ln \mathtt{P})/\mathtt{TSD}} \approx f(\tau_0) - \frac{\ln \mathtt{P}}{\mathtt{TSD}} f(\tau_0)(\tau-\tau_0) + \frac{1}{2}\left(\frac{\ln \mathtt{P}}{\mathtt{TSD}}\right)^2 f(\tau_0)(\tau-\tau_0)^2 \\
&= \mathtt{P}^{-\tau_0/\mathtt{TSD}} - \frac{\ln \mathtt{P}}{\mathtt{TSD}} \mathtt{P}^{-\tau_0/\mathtt{TSD}}(\tau-\tau_0) + \frac{1}{2}\left(\frac{\ln \mathtt{P}}{\mathtt{TSD}}\right)^2 \mathtt{P}^{-\tau_0/\mathtt{TSD}}(\tau-\tau_0)^2.
\end{aligned}$$

Substituting this approximation in (27) gives

$$\frac{1}{2}B\left(\frac{\ln \mathtt{P}}{\mathtt{TSD}}\right)^2 \mathtt{P}^{-\tau_0/\mathtt{TSD}}(\tau-\tau_0)^2 - \left(A + B\frac{\ln \mathtt{P}}{\mathtt{TSD}}\mathtt{P}^{-\tau_0/\mathtt{TSD}}\right)(\tau-\tau_0) + \left(A\tau_0 + B(\mathtt{P}^{-\tau_0/\mathtt{TSD}} - 1) - C\right) = 0.$$

Solving this quadratic for $\tau - \tau_0$ gives a very good approximation for τ. This is implemented in the Matlab code in Appendix B.4.

B Matlab code

B.1 packetSchedule()

This is the main function.

```
function [CN,T,t] = packetSchedule( SR, BCR, P, TSD, Q )
%[CN,T,t] = packetSchedule( SR_P, BCR_P, P, TSD, Q )
%  Determine the sequence of CNs for a single time slot.
%  This sequence is used in subsequent time slots by incrementing
%  modulo T all of the CNs that are different from T.
%
%  Defaults:  SR_P  = 100
%             BCR_P = 1
%             P     = 0.75
%             TSD   = 10
%             Q     = 30

if ~exist('SR','var'),
   SR = 100;
end
if ~exist('BCR','var'),
   BCR = 1;
end
if ~exist('P','var'),
   P = 0.75;
end
if ~exist('TSD','var'),
   TSD = 10;
```

```
end
if ~exist('Q','var'),
   Q = 30;
end

b = baseChannelTimes( BCR, P, TSD );
[w,N] = waveChannelTimes( SR, BCR, P, TSD );
T = N + Q;

times = [b,mod(N*TSD-w,TSD)];
CN = [T*ones(1,length(b)),floor(w/TSD)];
[t,j] = sort(times);
CN = CN(j);
```

B.2 baseChannelTimes()

This computes the transmission times for the base channel. It is subordinate to packetSchedule().

```
function [b,L] = baseChannelTimes( BCR_P, P, TSD )
% baseChannelTimes( BCR_P, P, TSD )
%  Compute a vector of packet transmission times for the base
%  channel for one time slot beginning at time 0.
%
%  Defaults:  BCR_P = 1
%             P     = 0.75
%             TSD   = 10

if ~exist('BCR_P','var'),
   BCR_P = 1;
end
if ~exist('P','var'),
   P = 0.75;
end
if ~exist('TSD','var'),
   TSD = 10;
end

L = ceil( -BCR_P*TSD*(1-P)/log(P) );
b = TSD * log( 1 + log(P)/BCR_P/TSD * (0:L-1) ) / log(P);
```

B.3 waveChannelTimes()

This computes the transmission times for a single wave. It is subordinate to packetSchedule().

```
function [w,N] = waveChannelTimes( SR, BCR, P, TSD, psi, phi, makePlot )
% waveChannelTimes( SR_P, BCR_P, P, TSD )
%  Compute a vector of packet transmission times for a single
%  wave that starts at time 0.
%
%  The result is accurate for any SR_P >= BCR_P.
%
%  Defaults:  SR_P  = 100
%             BCR_P = 1
%             P     = 0.75
%             TSD   = 10

if ~exist('SR','var'),
   SR = 100;
```

```
end
if ~exist('BCR','var'),
   BCR = 1;
end
if ~exist('P','var'),
   P = 0.75;
end
if ~exist('TSD','var'),
   TSD = 10;
end
if ~exist('psi','var'),
   psi = 1/4;
end
if ~exist('phi','var'),
   phi = 1/5;
end
if ~exist('makePlot','var')
   makePlot = 0;
end

mu = psi * (1-P) * (SR/BCR - 1);
N = ceil( log( 1 + (1/P)^(1-phi)*(1/P-1)*(SR/BCR - mu) )/log(1/P) ) - 1;
MWCR = (1-P)/(1-P^N)*(SR - mu*BCR);
tCrest = TSD * max([( N - log(MWCR/BCR)/log(1/P) ), 1]);
kappa = ( P^(-N+1) - 1 ) / ( (1/P) - 1 );

I0 = -BCR*TSD*(1-P)/log(P);
I1 = -BCR*TSD/log(P) * (P^(-N+tCrest/TSD) - 1);
I2 = (SR-mu*BCR)*(tCrest-TSD) + BCR*TSD/log(P) * kappa * (1-P^(tCrest/TSD-1));
I3 = SR*(2*TSD-tCrest) + BCR*TSD/log(P) * (P + kappa) * ((1/P)^(2-tCrest/TSD)-1);
I4 = mu*BCR*(tCrest-TSD);

L = ceil( I0 );
K1 = ceil( I0+I1 );
K2 = ceil( I0+I1+I2 );
K3 = ceil( I0+I1+I2+I3 );
K4 = round(TSD*SR);
K3 = min( [K3, K4] );

% part 4/4 of wave:
s1 = TSD * log( 1 - log(P)/BCR/TSD * ((L:(K1-1)) - I0) ) / log(1/P);

% part 3/4 of wave:
A = SR - mu*BCR;
B = -BCR * kappa * TSD/log(P) * P^(tCrest/TSD-1);
tau = solveIntegerCrossings( A, B, (K1:(K2-1)) - I0 - I1, P, TSD );
s2 = (N*TSD-tCrest)+tau;

% part 2/4 of wave:
A = SR;
B = -BCR * (P + kappa) * TSD/log(P);
tau = solveIntegerCrossings( A, B, (K2:(K3-1)) - I0 - I1 - I2, P, TSD );
s3 = (N-1)*TSD+tau;

% part 1/4 of wave:
tau = ((K3:(K4-1)) - I0 - I1 - I2 - I3)/(mu*BCR);
s4 = (N+1)*TSD-tCrest+tau;
```

```
w = [s1, s2, s3, s4];
if makePlot,
   figure
   plot( (L+1):K1, s1, 'b.' );
   hold on
   plot( (K1+1):K2, s2, 'r.' );
   plot( (K2+1):K3, s3, 'g.' );
   plot( (K3+1):K4, s4, 'm.' );
   title(['SR = ',num2str(SR),', BCR = ',num2str(BCR)])
end
```

B.4 solveIntegerCrossings()

This makes the calculation described in Appendix A. It is used by waveChannelTimes().

```
function tau = solveIntegerCrossings( A, B, C, P, TSD )
%tau = solveIntegerCrossings( A, B, C, P, TSD )
%
%  Find approximate solution to A*tau + B*(1-P^(-tau/TSD)) = C.
%
%  The normal usage is to pass a vector C of integer-spaced values,
%  in which case the equation is solved for each component of C.
%
%  Implemented with a quadratic Taylor expansion about tau0, where tau0
%  is the approximate solution obtained with a linear Taylor expansion.

if ~exist('P','var'),
   P = 0.75;
end
if ~exist('TSD','var'),
   TSD = 10;
end

tau0 = C/(A + B*log(P)/TSD);
ftau0 = P.^(-tau0/TSD);

a = 0.5*B*(log(P)/TSD)^2 * ftau0;
b = -(A + B*log(P)/TSD * ftau0);
c = C + B*(ftau0 - 1) - A*tau0;
zeta = (-b-sqrt(b.^2-4*a.*c) )./(2*a);
tau = tau0 + zeta;
```

Acknowledgments

Robert Chapman first observed the oscillatory behavior that indicated the deficiency of Design 2.

Originally, transmission times for the base and wave channels were computed independently. Armin Haken demonstrated the possibility of off-by-one errors in the number of wave channel packets produced. This was traced to the difference between I_0 and L and inspired the unified treatment given here.

References

[1] M. Luby, V. K Goyal, S. Skaria, and G. B. Horn. Wave and equation based rate control using multicast round trip time. *Comp. Comm. Rev.*, 32(4):191–214, October 2002. (This issue is Proc. ACM SIGCOMM 2002.).

[2] M. Luby and V. K Goyal. Wave and equation based rate control using multicast round trip time: Extended report. Technical Report DF2002-07-001, Digital Fountain, September 2002. Available on-line at `http://www.digitalfountain.com/technology/`.

[3] M. Luby, V. K Goyal, and S. Skaria. Wave and equation based rate control: A massively scalable receiver driven congestion control protocol. IETF Reliable Multicast Transport Working Group Internet-Draft, October 2001. Version 0, superseded by [4].

[4] M. Luby and V. K Goyal. Wave and equation based rate control building block. IETF Reliable Multicast Transport Working Group Internet-Draft, March 2002. Version 1, superseded by [5].

[5] M. Luby and V. K Goyal. Wave and equation based rate control building block. IETF Reliable Multicast Transport Working Group Internet-Draft, June 2002. Available on-line at `http://www.ietf.org/internet-drafts/draft-ietf-rmt-bb-webrc-02.txt`. Work in progress. Expires Dec. 2002.