Report Number:

2016-12

Multipath and Rate Stability

Junjie Liu and Roch A. Guérin

Originally Published In Proc. IEEE Globecom Conference - CQRM: Communication QoS, Reliability & Modeling Symposium

Multipaths and Rate Stability*

Junjie Liu and Roch Guérin
Department of Computer Science and Engineering
Washington University in St. Louis
St. Louis, Missouri 63130
Email: (junjie.liu,guerin)@wustl.edu

***Abstract*—Multipath solutions have been shown to help improve throughput, reliability and/or load balancing. This paper seeks to understand if and when they benefit *rate stability*. Rate stability is important to many real-time, interactive applications, *e.g.,* streaming video, but whether multipath solutions can help is unclear. Of relevance is the *time-scale* at which bandwidth changes are detected and acted upon to rebalance transmissions across paths. Consider two boundary cases: instantaneous detection and rate re-allocation, and a static rate assignment based on long-term path statistics. When transmissions can be instantaneously rebalanced across paths based on real-time link rate information, a multipath solution trivially improves rate stability (it all but eliminates rate variations). In contrast, when rate allocations are static, we find that multipath cannot improve upon the best single-path solution when buffers are large. When buffers are small (and coding is used to overcome losses), a multipath solution can, however, be beneficial even under a static rate allocation. The paper provides insight into when and how multipath solutions can help improve rate stability.**

I. INTRODUCTION

The control of end-to-end latency, be it in the form of average delay or jitter, is of importance to a wide range of applications and systems. This is true for real-time systems that often involve control loops with tight delay constraints [5], [18]. It is also of relevance to many audio and video applications [6], [12], [19], [23], where delay and jitter affect the size of playback buffers and the level of interactivity. Similarly, distributed multi-player games also exhibit sensitivity to both delay and jitter [11], [14], which determine the feasibility and accuracy of many game actions.

Delay itself is made up of propagation, transmission, and queueing delays, with the latter responsible for delay variations. Those variations arise from fluctuations in the transmission rate available to individual flows, *i.e.,* the bandwidth share they are able to get. In other words, delay variations are induced by variations in the end-to-end rate available to an individual flow. There are many causes for rate variations in packet networks. In wireless networks, phenomena such as fading, multipath propagation, shadowing, interferences, and mobility all affect transmission rates. Conversely, the high level of statistical multiplexing that is the norm in wired networks can give rise to significant fluctuations in the volume of traffic carried by an individual link. The resulting variations in link congestion in turn affect the transmission rates individual flows experience.

*This work was supported in part by NSF grant CNS-1361771.

Hence, developing solutions that can control rate variability, and therefore delay variability is of interest. In this paper, we explore whether a multipath approach can be effective in lowering end-to-end rate variability. Multipath is known to improve throughput and reliability in both wired and wireless networks [2], [9], [10], [15], [17], [30], and there is initial evidence that it could also benefit delay [20]. It is, therefore, natural to ask whether similar advantages extend to *rate stability*. A goal of this paper is to develop insight as to when and why this might be the case (or not). In exploring this question, because rate variations are temporal in nature, it is important to understand and specify the time-scale at which multipath decisions are made, and on the basis of what information. In other words, what is known about the transmission rates of individual paths, how fast do path rate changes become known to the source, and how is that information used to make decisions on which path(s) to use?

Specifically, we consider an environment where a source has access to multiple distinct paths to a destination, and is responsible for deciding how much traffic to forward on each path. Its goal is to realize a given (average) target rate at the receiver, while minimizing rate variations. Rates, and therefore rate variations, are measured at a time-scale that is application specific. For example, a real-time control process may be sensitive to rate changes at a time-scale below a millisecond or lower [13], while an interactive audio session or a robotic control loop might tolerate rate variations extending to a few tens of milliseconds [8], and an interactive game may be able to absorb rate variations over periods ranging from ten to over a hundred milliseconds [7].

In seeking to minimize rate variations, the sender decides (schedules) how much of its traffic to send on each path based on information available about individual path rates. This information and the associated scheduling decisions can be dynamic or static. In a dynamic scenario, the sender receives regular rate updates and uses them to change its scheduling decisions. In this case, the quality of the sender's decisions depends on the timeliness of the rate information it receives. Conversely, rate information at the sender can be static and in the form of path rate statistics measured over an extended period of time. In this case, the allocation of traffic across paths is itself static and computed based on those statistics.

For purpose of illustration, consider two (extreme) examples. Assume first a local network where path rates continuously monitored with measurements are instantaneously available at the sender, *i.e.,* there is no time-lag. In such a scenario, the sender also reacts immediately to rate changes and reallocates transmissions to maintain a steady rate[1] (assuming sufficient aggregate capacity across all paths). Under such assumptions, the availability of multiple paths can all but eliminate rate variability.

Consider next a scenario with a large time lag between changes in path rates, and when the sender becomes aware of them and reacts to reallocate transmissions. There are many possible causes for such latency, from large propagation delays, to low responsiveness of the rate monitoring mechanism itself, to high overhead in recomputing rate allocations in response to changes, etc. In such situations, a reactive scheme may be counter-productive, *i.e.,* implement rate reallocations that systematically trail path rate changes. Instead, a proactive approach that computes static rate allocations based on long-term path rate statistics is likely to be more effective. For example, consider a configuration involving a sender and a receiver connected by multiple paths, but separated by very large distances, so that the roundtrip propagation alone is over, say, 200 milliseconds. Assume further that rate variations occur close to the receiver on all paths. The sender is, therefore, aware of rate changes at best 100 milliseconds after they have occurred, so that dynamic rate updates are of limited benefit. Long-term rate statistics are then more useful, and can be used to compute how to distribute transmissions across paths to minimize rate variations at the receiver, *e.g.,* as proposed in [3]. In this high delay case, whether multiple paths improve rate variability over, say, using the best[2] path is unclear.

The paper explores this question through a stylized model that captures core aspects of a multipath solution, and makes the following contributions. It identifies that when the time lag between path rate changes and the sender's reaction to them is large so that a static rate allocation is in order, then a multipath solution needs not improve over the best single path solution. This is in contrast to the low delay case, where, as discussed above, a dynamic multipath rate allocation strategy outperforms the best single path solution. The paper also highlights the role of buffers in this outcome. When delays are large and buffers are small and, therefore, insufficient to avoid losses in the presence of rate fluctuations, a multipath solution coupled to application-level coding to recover from losses, can improve rate stability. However, when buffers are large enough to avoid data losses, the best single path solution always yields lower rate variability. The result holds even absent the impact of packet reordering across paths [28], and we offer evidence that the result also stands even when accounting for reordering.

The rest of the paper is organized as follow. Section II briefly reviews a number of related works. Section III introduces the model used to capture the performance of a multipath solution in terms of rate stability. Different configurations are considered to isolate the impact of delay, buffers, and reordering at the receiver. Section IV presents our solution method and its results, including when and why a single path solution outperforms a multipath solution. Finally, Section V summarizes the paper's findings and identifies a number of possible extensions.

II. Related Works

As indicated earlier, the topic of multipath has been of interest for a number of years. A comprehensive survey is beyond our scope. Instead, we sample works representative of the benefits of multipaths, and review briefly a paper with motivations similar to ours, namely, improving rate stability.

Because multipath solutions offer access to more plentiful and diverse network resources, they are natural candidates for improving either reliability, or throughput, or both. [9] investigated the use of multipaths to better support high bit rate applications in low-bandwidth networks, while [10] sought to increase aggregate bandwidth in inter-domain routing. Similarly motivated investigations were carried out for sensor networks, where link bandwidth is low and highly variable, *e.g.,* see [30] for a recent survey. The benefits of path diversity for video streaming was explored in [2] and [17]. Both papers highlighted increases in throughput and reduction in loss correlation, while [2] also demonstrated improvements in video quality when used in combination with multiple description coding. The use of coding to improve throughput was further investigated in [31], [32]. Additionally, the use of multipath to improve end-to-end reliability and/or fault tolerance was studied in a number of earlier works, *e.g.,* [4], [15], [21], [22], [24], with [29] offering the first theoretical investigation of the problem.

Surprisingly, the use of multipath to improve rate stability has received little attention to-date. One exception is [3], which shares similar goals as this paper. Specifically, [3] introduces a distributed optimization framework to compute multipath solutions that minimize rate variance while meeting minimum average rate guarantees. For analytical tractability, [3] minimizes the sum of individual link rate variances across a multipath rather than the end-to-end path rate variance, as we do in this paper. In addition, [3]'s focus is on a protocol for realizing a solution, rather than understanding when and why a multipath solution can lower rate variability. Another related work is [27], which focused on live streaming applications and investigated the use of multipaths to reduce jitter[3]. It relied on duplicate transmissions over two paths with nearly equal delays, and demonstrated by simulation improvements in jitter over the best single path solution. The paper, however, did not seek to elucidate when and why multipath transmissions would in general reduce jitter and/or rate variability.

III. Model Description and Problem Statement

This section introduces our model for comparing multipath and single-path solutions in terms of rate stability. The problem is first presented in a general form, from which a simplified yet representative model is extracted for analytical tractability.

[1]The CONGA system [1] offers a close approximation of such a behavior.
[2]See Section III for a more precise definition of "best."

[3]As discussed in the next section, jitter and rate variations are closely related.

A. A General Multipath Model

A network is represented by a directed graph $G(V, E)$. Vertices (V) correspond to network/terminal nodes and edges (E) to links between nodes. For vertices $i, j \in V, (i, i') \in E$ denotes the edge connecting vertex i to vertex i'. A path through the network is a sequence of links of the form $p \equiv \{(s, i_1); (i_1, i_2); \ldots; (i_m, d)\}$, where s denotes the *source* node and d is the *destination* node. The length of path p, in number of links, is $m + 1$. Consider the configuration of

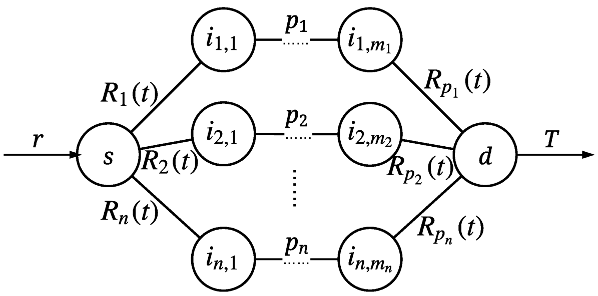

Fig. 1: A general multipath configuration.

Fig. 1, consisting of source node s and destination nodes d connected by a set of n paths $P = \{p_1, p_2, \ldots, p_n\}$, where $p_j \equiv \{(s, i_{j,1}); (i_{j,1}, i_{j,2}), \ldots, (i_{j,m_j}, d)\}$ is the $j^{\text{th}}, 1 \le j \le n$, path in P. Node s is sourcing a flow of rate of r bits/sec, and is making scheduling decisions that determine its transmission rate[4] on each path. Specifically, $R_j(t), 1 \le j \le n$, denotes the flow's transmission rate on path j at time t.

Transmissions proceed in a store-and-forward fashion on each path, with buffers available at nodes to store bits whenever the incoming rate is higher than the outgoing rate. In particular, a flow's incoming and outgoing rates at each node are assumed to be time-varying. As mentioned earlier, there are multiple possible causes for such variations, *e.g.,* fluctuations in link capacity in wireless networks or the impact of statistical multiplexing in wired networks. Of most interest are the flow's outgoing rates from the penultimate nodes on each path, *i.e.,* $R_{p_j}(t)$ for path j as shown in Fig. 1, as together they determine the aggregate incoming flow rate at the destination d.

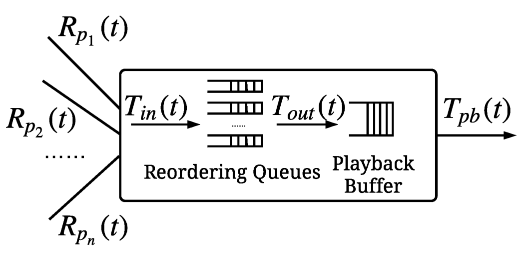

Fig. 2: Destination node structure

The structure of the destination node d is expanded in Fig. 2, which shows the incoming rates $R_{p_j}(t), 1 \le j \le n$, for each of the paths between s and d. Together they make up the total raw incoming rate at the destination node, $T_{in}(t) = \sum_{j=1}^{n} R_{p_j}(t)$. However, because bits sent over different paths need not arrive in order, they are first fed to *reordering queues*. Bits are read from the reordering queues as soon as they can be forwarded in order to the *playback buffer*. Denote as $\widehat{R}_{p_j}(t), 1 \le j \le n$, the rate at which bits exit the reordering queues of each path, so that $T_{out}(t) = \sum_{j=1}^{n} \widehat{R}_{p_j}(t)$ is the incoming rate to the playback buffer. Note that the $\widehat{R}_{p_j}(t)$'s have dependencies on each other and on the $R_{p_j}(t)$'s, since the arrival of a bit on one path can unlock the delivery of bits from other paths.

The rate into the playback buffer, $T_{out}(t)$, is the quantity of interest in that it determines the extent to which the playback rate, $T_{pb}(t)$, differs from the application rate, r. Some applications rely on an initial *buffering* phase to absorb fluctuations in $T_{out}(t)$. Buffering, however, introduces additional playback delay that can affect application performance, so that keeping it small is desirable. This is what motivates our goal of keeping variations in $T_{out}(t)$ small (the smaller they are, the smaller the playback buffer needs to be to maintain a constant playback rate of r). Specifically, we seek a solution to the following optimization problem $\boldsymbol{P}$:

$$\begin{aligned} \boldsymbol{P} : \forall t, \vec{R}^*(t) &= \arg\min_{\vec{R}(t)} \operatorname{Var}\left(T_{out}(t)\right), \\ s.t.\ E[T_{out}(t)] &\ge r \end{aligned} \tag{1}$$

where $\vec{R}(t) = [R_1(t), R_2(t), \ldots, R_n(t)]$ is the vector of transmission rates on paths $p_1, p_2, \ldots, p_n$, at the source node s, and r is the original source rate. Our goal is to minimize the variance of the rate $T_{out}(t)$ at the destination node d, while preserving the original source rate r.

Before turning to problem $\boldsymbol{P}$, we briefly discuss the close relationship that exists between rate variability and jitter. Jitter measures variations in the time between consecutive packets (bits). This time is inversely proportional to the transmission rate, with variations in rate contributing directly to jitter. Defining jitter as $J = \operatorname{Var}\left(\frac{1}{T_{out}(t)}\right)$ and denoting $T_{out}(t) = r + x(t)$ with $E[x(t)] = 0$, it is easy to show (using a simple Taylor series expansion) that for $\frac{x(t)}{r}$ small

$$J \sim \frac{\operatorname{Var}\left(T_{out}(t)\right)}{r^4}$$

In other words, minimizing $\operatorname{Var}\left(T_{out}(t)\right)$ minimizes jitter.

Solving problem $\boldsymbol{P}$ in its most general instance is complex, especially when path rates exhibit dependencies. We, therefore, introduce several restrictions to ensure tractability, while still capturing key features of multipath solutions. We first assume that paths are independent and that rate variations are contributed only by the last link of a path, so that a path can be nominally viewed as consisting of only two links as shown in Fig. 3 for a two paths scenario. Rate variations on a path are also streamlined and limited to an ON-OFF pattern, *i.e.,* akin to a Gilbert-Elliot channel [16]. Specifically, path p_i is in the ON state with probability (fraction of time) ρ_i during which

[4]For the sake of simplicity, we assume continuous transmissions and ignore the discretization imposed by packet boundaries.

it offers a transmission rate $C_j > r$ to the source[5]. It is in the OFF state the rest of the time with a transmission rate of 0. Each path is further assumed to have an average rate $E[R_{p_j}(t)]$ that exceeds r, so that any single path can fully accommodate the source rate. In other words, $\forall j$, $E[R_{p_j}(t)] = \rho_j C_j \geq r$. Note also that $\mathrm{Var}(R_{p_j}(t)) = \rho_j(1-\rho_j)C_j^2$.

In spite of these simplified assumptions, a wide range of path rate variations are still feasible, and the focus is on their impact on the variability of $T_{out}(t)$ without concerns for constraints such as bandwidth limitations on individual paths. We further consider two special (extreme) cases of the

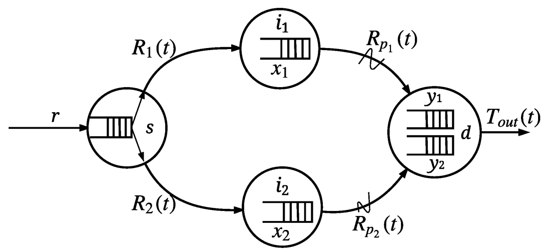

Fig. 3: A basic two-path instance of problem $\boldsymbol{P}$

simplified configuration of Fig. 3, through which we seek to explore if and how a multipath solution can be useful.

(1) Delays (from path rate change to sender rate reallocation) are negligible, so that the source instantaneously reacts to rate changes, and dynamically adjusts its transmission decisions to minimize rate variations at the destination.

(2) Delays are large, so that the source "never" learns of or reacts to rate changes[6]. It is only aware of the (long-term) rate statistics of each path, and makes static rate allocation decisions to minimize $\mathrm{Var}(T_{out})$ based solely on that information.

This latter scenario is further split into two sub-cases.

(2a) Network buffers (at nodes i_1 and i_2 of Fig. 3), are small. As the source does not adapt to rate variations, this implies data loss. To counter the effect of those losses, we assume the use of perfect codes, *e.g.,* Fountain codes [26], at the source[7].

(2b) Network buffers are large enough to avoid all network losses. For simplicity, we ignore the possibility of time-outs, so that retransmissions are not needed, and assume similar propagation delays on all paths (data still arrives out of order when one or more paths is in the OFF state). Hence, while reordering was absent from the other scenarios, it now induces variations in $T_{out}(t)$ and its impact needs to be accounted for.

IV. Solution Method

This section presents our approach to solving problem $\boldsymbol{P}$ in the two configurations introduced in the previous section. We start with configuration 1, where the source is instantaneously aware of rate changes on each path and can adapt its transmissions accordingly.

A. Scenario 1: Instantaneous Source Adaptation

For ease of comparison, we first describe the solution to problem $\boldsymbol{P}$ in a single path scenario. With just one path and short transmission delays, the transmission rate at the source, $R(t)$, is essentially the rate seen by the destination[8], *i.e.,* $R(t) = T_{out}(t)$. In this case, the source's only decision is how high to set $R(t)$ after a period of interruption on the path to the destination. When the path comes back ON, the source has accumulated data in its buffer, there is then a trade-off between how fast it empties the buffer and the magnitude of the variations in T_{out} it induces. This can be specified through the following statistics

$$T_{out}(t) = R(t) = \begin{cases} r', & \text{w.p. } p_1 \\ r, & \text{w.p. } p_2 \\ 0, & \text{w.p. } 1-\rho\,, \end{cases}$$

where $r' \leq C$ is the buffer draining rate at the source, and $1-\rho$ is the probability that the path is in the OFF state.

Our goal is to minimize $\mathrm{Var}(T_{out}(t)) = r'r(1-\rho)$ subject to the constraint $E[R(t)] = p_1 r' + p_2 r = r$. Since $\mathrm{Var}(T_{out}(t))$ is monotonically increasing in r', the solution is to pick the smallest possible value for r' that satisfies the constraint $E[R(t)] = r$. This constraint together with the fact that $p_1 + p_2 = \rho$ gives

$$p_1 = \frac{(1-\rho)r}{r'-r} \quad \text{and} \quad p_2 = \frac{\rho r' - r}{r'-r}$$

which immediately implies $r' = \frac{r}{\rho}$, and therefore the following solution under the assumption of a single path:

$$\mathrm{Var}(T^*_{out,1}(t)) = \frac{r^2(1-\rho)}{\rho} \tag{2}$$

Eq. (2) is monotonically increasing in ρ, which is intuitive, *i.e.,* a more reliable path results in lower rate variance.

Consider now the case of two paths. The source buffers data when both paths are in the OFF state. When its buffer is not empty and either path is in the ON state, the source transmits at rate r' (typically using only one path). Otherwise, it transmits at rate r. This gives rise to a similar formulation as in the single path case, with the modification that the probability that the source's rate is 0 is now equal to $p_3 = (1-\rho_1)(1-\rho_2)$, where $\rho_j, j = 1, 2$, is the probability that path j is in the ON state.

Using again $E[R(t)] = r$, we now get

$$p_1 = \frac{(1-a)r}{r'-r} \quad \text{and} \quad p_2 = \frac{ar' - r}{r'-r}$$

where $a = (1-\rho_1)(1-\rho_2)$. This also gives $\mathrm{Var}(T_{out}(t)) = rr'(1-\rho_1)(1-\rho_2)$, and minimizing $\mathrm{Var}(T_{out}(t))$ still calls for choosing the smallest possible value for r' while ensuring $E[R(t)] = r$. This gives $r' = \frac{r}{1-(1-\rho_1)(1-\rho_2)}$, and therefore

$$\mathrm{Var}(T^*_{out,2}(t)) = \frac{r^2(1-\rho_1)(1-\rho_2)}{1-(1-\rho_1)(1-\rho_2)} \tag{3}$$

[5] C_j is the link bandwidth available to the source after accounting for the impact of other traffic.

[6] Or at least, it learns about or reacts to them too late to meaningfully affect rate variations at the destination.

[7] Note that this comes at the cost of an increase in the source transmission rate because of coding overhead.

[8] We assume that to avoid losses the source stops transmitting, *i.e.,* it buffers data, as soon as it detects that the last link on its path is in the OFF state.

Comparing Eqs. (2) and (3) readily gives that unless $\rho_j = 1$ for either one of the two paths, a two paths solution can reduce rate variability. Note that the solution is easily generalized to n paths, with the following variance for the best rate setting:

$$\mathrm{Var}(T^*_{out,n}(t)) = \frac{r^2 \prod_{j=1}^{n}(1-\rho_j)}{1-\prod_{j=1}^{n}(1-\rho_j)} \tag{4}$$

where ρ_j is the probability that the i^{th} path is in the ON state. The result is summarized in the following proposition.

Proposition 1: When paths exhibit rate variations that are instantaneously detectable at the source so that it can immediately react to them, a multipath solution can lower rate variability at the destination. In addition, the ability to lower rate variability increases with the number of paths available. The benefits of a multipath solution articulated in Proposition 1 are illustrated in Fig. 4 for different number of paths and path statistics (values of ρ), under the assumption of identical paths.

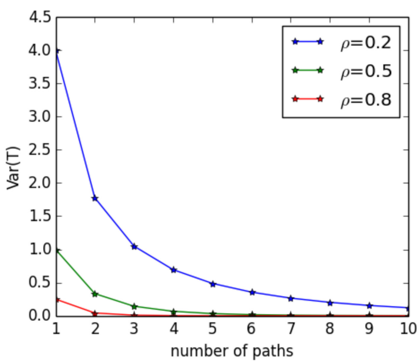

Fig. 4: Instantaneous source adaptation: Rate variance as a function of number of paths and path statistics

B. Scenario 2: Static Rate Allocation

In this scenario, the time lag between the occurrence of path rate changes and when the source reacts to them, is too large for the rate adaptation of the source to mitigate rate variations at the destination. Instead, the rate statistics of each path, *i.e.,* the mean and variance of the end-to-end rate, are used by the source to make *static* rate allocation decisions that minimize $\mathrm{Var}(T_{out}(t))$. Two sub-cases are considered: A small buffer scenario where coding is used to overcome losses, and a scenario where network buffers are large enough to avoid losses caused by path rate variations.

1) Scenario 2a: Small Buffers: Let $\widehat{r} \geq r$ denote the source transmission rate accounting for coding overhead, where successful decoding at the destination requires[9] $E[T_{out}(t)] > r$. The transmission rate on path j is then of the form:

$$R_j(t) = R_{p_j}(t) = \begin{cases} \alpha_j \widehat{r}, & \text{w.p. } \rho_j \\ 0, & \text{w.p. } 1-\rho_j \end{cases}$$

where α_j is the (static) fraction of the source rate plus coding that is sent on path j, and ρ_j is the fraction of time that path j is in the ON state. The source is assumed to know all the ρ_j's and to use them to compute α_j values that minimize rate variations at the destination. Specifically, the combination of small buffers and static rate allocations implies that

$$T_{out}(t) = \sum_j \alpha_j \delta_j(t) \widehat{r}$$

where $\delta_j(t) = 1$, if the state of path j is ON and 0 otherwise. This yields the following mean and variance for $T_{out}(t)$

$$E[T_{out}(t)] = \sum_j \alpha_j \rho_j \widehat{r} \tag{5}$$

$$\mathrm{Var}(T_{out}(t)) = \widehat{r}^2 \sum_j \alpha_j^2 \rho_j (1-\rho_j) \tag{6}$$

Note that Eq. (5) implies that the smallest coding rate that allows successful decoding at the destination is

$$\widehat{r}_{\min} \approx \frac{r}{\sum_j \alpha_j \rho_j} \tag{7}$$

Choosing $\widehat{r} = \widehat{r}_{\min}$, therefore, lowers $\mathrm{Var}(T_{out}(t))$ to

$$\mathrm{Var}(T_{out}(t)) = \frac{r^2 \sum_j \alpha_j^2 \rho_j (1-\rho_j)}{\left(\sum_j \alpha_j \rho_j\right)^2} \tag{8}$$

When coding is used, the destination experiences rate variations whenever decoding stalls, *i.e.,* when $T_{out}(t) < r$. Minimizing rate variations, therefore, calls for minimizing $\mathrm{Prob}(T_{out}(t) < r)$. From Chebyshev's inequality this is reasonably approximated by minimizing $\mathrm{Var}(T_{out}(t))$. Solving Eq. (8) for the values of α_j that minimize $\mathrm{Var}(T_{out}(t))$ readily yields the following proposition.

Proposition 2: When paths exhibit rate variability, *i.e.,* $\rho_j \neq 1$, a static rate allocation that uses coding to eliminate losses, minimizes rate variability at the destination by splitting traffic across paths as follows

$$\alpha_j = \frac{\prod_{k \neq j}(1-\rho_k)}{\sum_m \prod_{k \neq m}(1-\rho_k)} \tag{9}$$

The benefits of a multipath solution are illustrated in Fig. 5 for the case of two paths (assuming $C_j \gg r$ so that $\rho_j \approx 0$ is feasible while still satisfying $\rho_j C_j > r$). It shows the reduction in rate variance of the best multipath solution over the best single-path solution. The figure also highlights that this reduction is maximum (50%) when $\rho_1 = \rho_2$.

2) Scenario 2b: Large Buffers: Buffers are assumed large enough to absorb link rate fluctuations and, therefore, eliminate network losses. Rate variations at the destination now arise because of both path rate variations, *i.e.,* as captured by $T_{in}(t)$, and data reordering, *i.e.,* as reflected in $T_{out}(t)$ (see Fig. 2). To better assess the significance of each factor, we first investigate the impact of path rate variations before accounting for the additional effect of reordering.

The introduction of network buffers adds complexity as buffer contents must now be tracked along with path states.

[9]This can imply relatively large encoding and decoding delays, but we ignore this aspect for the time being.

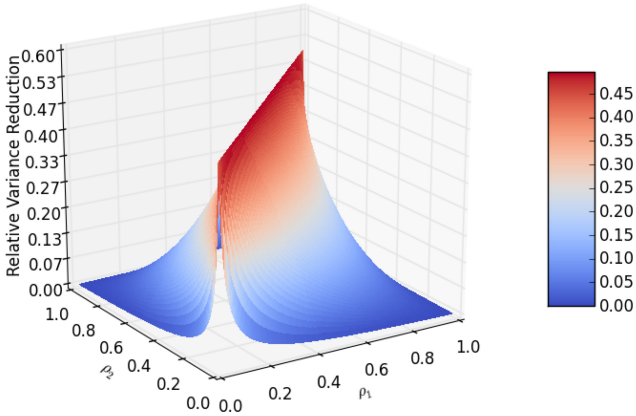

Fig. 5: Relative multipath variance reduction

Specifically, following the notation of Fig. 3, the rate $R_{p_j}(t)$ contributed by path j at the destination can take three values: C_j (link is ON, buffer is not empty); $R_{p_j}(t) = \alpha_j r$ (link is ON, buffer is empty); and 0 (link is OFF), where as before α_j is the fraction of the source's transmissions sent on path j. Computing $\mathrm{Var}(T_{in}(t)) = \mathrm{Var}\left(\sum_j R_{p_j}(t)\right)$ calls for characterizing the odds of each rate value on each path. A standard analysis yields

$$R_{p_j}(t) = \begin{cases} C_j, & \text{w.p. } \rho_j - F_j^{(1)}(0) \\ \alpha_j r, & \text{w.p. } F_j^{(1)}(0) \\ 0, & \text{w.p. } 1-\rho_j \end{cases} \tag{10}$$

where $F_j^{(1)}(0) = \frac{\rho_j C_j - \alpha_j r}{C_j - \alpha_j r}$ is the probability that the link buffer is empty and the link is ON.

Since large buffers eliminate losses, $E[R_{p_j}(t)] = \alpha_j r$, and Eq. (10) gives $\mathrm{Var}(R_{p_j}(t)) = (1-\rho_j)C_j\alpha_j r$. Hence,

$$E[T_{in}(t)] = r \tag{11}$$

$$\mathrm{Var}(T_{in}(t)) = \sum_j C_j \alpha_j r(1-\rho_j) \tag{12}$$

where Eq. (12) relies on our assumption of independent paths. The result implies that under a static allocation, in the absence of losses (and without coding) *no multipath solution can improve* $\mathrm{Var}(T_{in}(t))$ *over the best single-path solution.* Specifically, denoting $d_j = (1-\rho_j)C_j$ as the *rate deficit* of path j, we have more formally

Proposition 3: Under a static rate allocation and in the absence of losses, and therefore coding, let $j^* = \arg\min_j (1-\rho_j)C_j$, then $\mathrm{Var}(T_{in}(t)) \geq C_{j^*}(1-\rho_{j^*})r$, *i.e.,* the variance of the received rate before data reordering is minimized by assigning all the traffic to the path with the lowest rate deficit.

Proposition 3 establishes that with large buffers and delays high enough that dynamic rate adaptation is not feasible, a multipath solution does not reduce variations in the incoming network rate, $T_{in}(t)$. Next, we turn to assessing if this also holds for $T_{out}(t)$ that accounts for the impact of reordering. We focus on the simplest scenario, namely, all paths have the same propagation delay so that rate variations on individual paths are the only cause for out-of-order data.

Intuitively, adding a reordering constraint should only make matters worse. Analyzing this configuration is, however, complex because ordering couples rates across paths, *i.e.,* individual paths cannot anymore be studied in isolation. Furthermore, variations in $T_{out}(t)$ and $T_{in}(t)$ differ in both range and statistics. Accounting for these different factors makes a direct analysis challenging, but as we state next, two observations render it tractable (though not simple).

Proposition 4: Let x_j and y_j denote the content of the link and reordering buffers of path j, respectively. For simplicity, we consider the case of only two paths. We have

1) At any time, $(y_1 = 0) \vee (y_2 = 0)$, *i.e.,* at least one of the two reordering buffers is empty.
2) If $x_1 \geq \frac{\alpha}{1-\alpha}x_2$, then $y_1 = 0$, and if $x_1 \leq \frac{\alpha}{1-\alpha}x_2$, then $y_2 = 0$, where as before α is the fraction of traffic sent on path 1.

The proof can be found in [25].

Using Proposition 4, we derive the distribution of $T_{out}(t)$.

$$T_{out}(t) = \begin{cases} \frac{C_1}{\alpha}, & \text{w.p. } P_{10}(x_2 < \frac{1-\alpha}{\alpha}x_1) + P_{11}(x_2 < \frac{1-\alpha}{\alpha}x_1) \\ \frac{C_2}{1-\alpha}, & \text{w.p. } P_{01}(x_2 > \frac{1-\alpha}{\alpha}x_1) + P_{11}(x_2 > \frac{1-\alpha}{\alpha}x_1) \\ r, & \text{w.p. } F_1^{(1)}(0) \cdot F_2^{(1)}(0) \\ 0, & \text{w.p. } F_1^{(0)}(\infty) \cdot F_2^{(0)}(\infty) + P_{10}(x_2 \geq \frac{1-\alpha}{\alpha}x_1) \\ & \quad +P_{01}(x_2 \leq \frac{1-\alpha}{\alpha}x_1) \end{cases} \tag{13}$$

where $P_{kl}(A), k, l \in \{0,1\}$ is the probability of the event A while the link of path 1 is in state k and the link of path 2 is in state l, and $F_j^{(k)}(x), j \in \{1,2\}, k \in \{0,1\}$, is the probability that the content of the link buffer of path j is less than equal to x and the link is in state k. The derivation of Eq. (13) and expressions for $P_{kl}(A)$ and $F_j^{(k)}(x)$ are again in [25].

With the distribution of $T_{out}(t)$ in hand, it is possible to obtain an expression for $\mathrm{Var}(T_{out}(t))$. Its complexity, however, offers little analytical insight, and we instead offer a representative example of its behavior in Fig. 6. The figure displays the value of $\mathrm{Var}(T_{out}(t))$ as the rate allocation parameter α varies from 0 to 1. It illustrates that allocating the entire flow to the path with the smaller rate deficit again yields the lowest rate variance. As a matter of fact, in this example, both single-path solutions outperform all multipath options. This illustrates the added penalty that reordering imposes as soon as a second path is used, even if only rarely, *i.e.,* data is held-up whenever one path falls behind. The next proposition formalizes the result, where we again use a two paths scenario for ease of exposition.

Proposition 5: Under a static rate allocation and in the absence of losses, and therefore coding, let $j^* = \arg\min_j (1-\rho_j)C_j$, then $\mathrm{Var}(T_{out}(t)) \geq C_{j^*}(1-\rho_{j^*})r$, *i.e.,* the variance of the received rate after data reordering is minimized by assigning all the traffic to the path with the lowest rate deficit. The proof can again be found in [25].

V. Conclusion & Future Work

The paper demonstrates that multipath solutions need not always reduce rate variability. The outcome depends on the

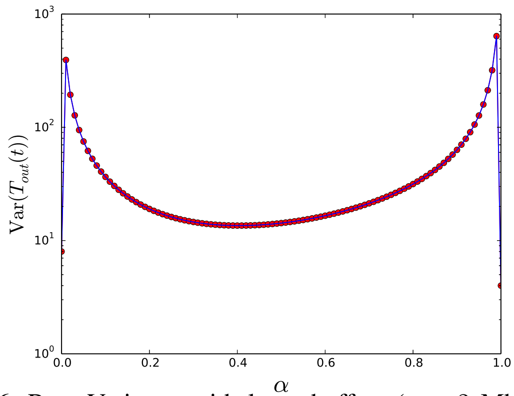

Fig. 6: Rate Variance with large buffers ($r = 2$ Mb/s, $C_1 = 4$ Mb/s, $C_2 = 8$ Mb/s, both paths have the same average time 100ms in ON and OFF state , *i.e.,* $\rho_1 = \rho_2 = \frac{1}{2}$).

extent to which dynamic rate adaption is possible. When latency is minimal so that instantaneous rate adaptation is feasible, a mutipath solution reduces rate variability. In contrast, when latency is so high that dynamic rate adaptation is impractical and only a static rate allocation is meaningful, then a multipath solution does not lower rate variability; at least not unless coding is used to overcome losses. Hence, there exists a *"crossover"* threshold in the latency with which rate adaptation can be performed. When latency exceeds this threshold, so that reactions to rate changes across paths are very slow, a multipath solution is unable to reduce rate variations. In contrast, when latency is below the threshold so that the source can detect and adapt to rate changes early enough, then a multipath solution can lower rate variability.

Exploring this question further and characterizing this threshold is of interest, as is accounting for differences in propagation delays across paths and the impact of retransmissions because of losses or time-outs. These are topics we expect to investigate in the future.

REFERENCES

[1] M. Alizadeh, T. Edsall, S. Dharmapurikar, R. Vaidyanathan, K. Chu, A. Fingerhut, V. T. Lam, F. Matus, R. Pan, N. Yadav, and G. Varghese. CONGA: distributed congestion-aware load balancing for datacenters. In *Proc. ACM SIGCOMM'14*, Chicago, IL, August 2014.

[2] J. Apostolopoulos and M. Trott. Path diversity for enhanced media streaming. *IEEE Comm. Mag.*, 42(8), August 2004.

[3] B. Arzani, R. Guerin, and A. Ribeiro. A distributed routing protocol for predictable rates in wireless mesh networks. In *Proc. IEEE ICNP*, Austin, TX, October 2012.

[4] T. Babbitt, C. Morrell, and B. Szymanski. Self-selecting reliable path routing in diverse wireless sensor network environments. In *Proc. IEEE ISCC*, Sousse, Tunisia, July 2009.

[5] K. P. Birman. *Reliable distributed systems – Technologies, web services and applications*. Springer, 2005.

[6] G. Blakowski and R. Steinmetz. A media synchronization survey: Reference model, specification, and case studies. *IEEE J. Select. Areas Comm.*, 14(1), January 1996.

[7] T. Burger. How fast is realtime? human perception and technology. PubNub, February 2015. See https://www.pubnub.com/blog/2015-02-09-how-fast-is-realtime-human-perception-and-technology/.

[8] G. Buttazzo. HARTIK: a real-time kernel for robotics applications. In *Proc. IEEE RTSS'93*, Raleigh-Durham, NC, December 1993.

[9] J. Chen, S. Chan, and V. Li. Multipath routing for video delivery over bandwidth-limited networks. *IEEE J. Select. Areas Commun*, 22(10), 2004.

[10] X. Chen, M. Chamania, A. Jukan., A. Drummond, and N. D. Fonseca. On the benefits of multipath routing for distributed data-intensive applications with high bandwidth requirements and multidomain reach. In *Proc. Comm. Netw. & Serv. Research Conf. (CNSR'09)*, Moncton, NB, May 2009.

[11] M. Claypool and K. Claypool. Latency and player actions in online games. *Commun. ACM*, 49(11):40–45, November 2006.

[12] M. Claypool and J. Tanner. The effects of jitter on the peceptual quality of video. In *Proc. ACM MULTIMEDIA'99*, pages 115–118, Orlando, Florida, USA, 1999.

[13] T. Cucinotta, A. Mancina, G. Anastasi, G. Lipari, L. Mangeruca, R. Checcozzo, and F. Rusina. A real-time service-oriented architecture for industrial automation. *IEEE Trans. Indus. Informatics*, 5(3), 2009.

[14] M. Dick, O. Wellnitz, and L. Wolf. Analysis of factors affecting players' performance and perception in multiplayer games. In *Proc. NetGames'05*, Hawthorne, NY, October 2005.

[15] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly-resilient, energy-efficient multipath routing in wireless sensor networks. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(4), 2001.

[16] E. Gilbert. Capacity of a burst-noise channel. *Bell System Technical Journal*, 39(5), September 1960.

[17] L. Golubchik, J. Lui, T. Tung, A. Chow, W.-J. Lee, G. Franceschinis, and C. Anglano. Multi-path continuous media streaming: what are the benefits? *Performance Evaluation*, 49(1-4), September 2002.

[18] G. W. Irwin, editor. *The engineering of complex real-time computer control systems*, volume 11(3) of *Special Issue of Real-Time Systems*. Kluwer Academic Publishers, 1996.

[19] Y. Ishibashi and S. Tasaka. A comparative survey of synchronization algorithms for continuous media in network environments. In *Proc. IEEE LCN'00*, Tampa, FL, November 2000.

[20] U. Javed, M. Suchara, J. he, and J. Rexford. Multipath protocol for delay-sensitive traffic. In *Proc. COMSNETS'09*, Bangalore, India, January 2009.

[21] A. Khandani, J. Abounadi, E. Modiano, and L. Zheng. Reliability and route diversity in wireless networks. *IEEE Trans. Wireless Comm.*, 7(12), December 2008.

[22] T. Korkmaz and K. Saraç. Characterizing link and path reliability in large-scale wireless sensor networks. In *Proc. IEEE WiMob*, Niagara Falls, Canada, October 2010.

[23] N. Laoutaris and I. Stavrakakis. Intrastream synchronization for continuous media streams: a survey of playout schedulers. *IEEE Network*, 16(3), May 2002.

[24] S.-J. Lee and M. Gerla. Split multipath routing with maximally disjoint paths in ad hoc networks. In *Proc. IEEE ICC*, Beijing, China, May 2001.

[25] J. Liu and R. Guérin. Multipath and rate stability. Tech. rep., Washington University in St. Louis, 2015. Available at http://students.cec.wustl.edu/~junjie.liu/multipath_tech_report.pdf.

[26] D. MacKay. Fountain codes. *IEE Proceedings – Communications*, 152(6), December 2005.

[27] T. Okuyama, K. Yasukawa, and K. Yamaoka. Nearly equal delay path set configuration (NEED-PC) for multipath delay jitter reduction. *IEICE Trans. Commun.*, 91(3), 2008.

[28] V. Paxson. End-to-end routing behavior in the Internet. In *Proc. ACM SIGCOMM*, Stanford, CA, August 1996.

[29] M. Rabin. Efficient dispersal of information for security, load-balancing, and fault tolerance. *Journal of the ACM*, 32(2), April 1989.

[30] M. Radi, B. Dezfouli, K. A. Bakar, and M. Lee. Multipath routing in wireless sensor networks: Survey and research challenges. *Sensors*, 12(1), January 2012.

[31] A. Tsirigos and Z. Haas. Analysis of multipath routing, part 1: The effect on the packet delivery ratio. *IEEE Trans. Wireless Comm.*, 3(1), January 2004.

[32] E. Vergetis, R. Guérin, and S. Sarkar. Realizing the benefits of user-level channel diversity. *ACM Computer Communication Review*, 35(5), October 2005.