

An axiomatic perspective on the performance effects of end-host path selection

Journal Article**Author(s):**

Scherrer, Simon; [Legner, Markus](#) ; Perrig, Adrian; Schmid, Stefan

Publication date:

2021-11

Permanent link:

<https://doi.org/10.3929/ethz-b-000510875>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

Performance evaluation 151, <https://doi.org/10.1016/j.peva.2021.102233>

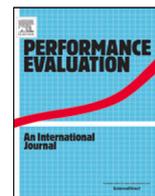
Funding acknowledgement:

182005 - ESCALATE: Efficient and Scalable Algorithms for Large Flow Detection (SNF)



Contents lists available at ScienceDirect

Performance Evaluation

journal homepage: www.elsevier.com/locate/peva

An axiomatic perspective on the performance effects of end-host path selection



Simon Scherrer^{a,*}, Markus Legner^a, Adrian Perrig^a, Stefan Schmid^{b,c,d}

^a Department of Computer Science, ETH Zurich, Zurich, Switzerland

^b Faculty of Electrical Engineering and Computer Science, TU Berlin, Berlin, Germany

^c Fraunhofer SIT, Darmstadt, Germany

^d Faculty of Computer Science, University of Vienna, Vienna, Austria

ARTICLE INFO

Article history:

Available online 28 September 2021

Keywords:

Computer networks
Congestion control
Path selection

ABSTRACT

In various contexts of networking research, end-host path selection has recently regained momentum as a design principle. While such path selection has the potential to increase performance and security of networks, there is a prominent concern that it could also lead to network instability (i.e., flow-volume oscillation) if paths are selected in a greedy, load-adaptive fashion. However, the extent and the impact vectors of instability caused by path selection are rarely concretized or quantified, which is essential to discuss the merits and drawbacks of end-host path selection.

In this work, we investigate the effect of end-host path selection on various metrics of networks both qualitatively and quantitatively. To achieve general and fundamental insights, we leverage the recently introduced axiomatic perspective on congestion control and adapt it to accommodate joint algorithms for path selection and congestion control, i.e., *multi-path congestion-control* protocols. Using this approach, we identify equilibria of the multi-path congestion-control dynamics and analytically characterize these equilibria with respect to important metrics of interest in networks (the “axioms”) such as efficiency, fairness, and loss avoidance. Moreover, we analyze how these axiomatic ratings for a general network change compared to a scenario without path selection, thereby obtaining an interpretable and quantitative formalization of the performance impact of end-host path-selection. Finally, we show that there is a fundamental trade-off in multi-path congestion-control protocol design between efficiency, stability, and loss avoidance on one side and fairness and responsiveness on the other side.

© 2021 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Path selection performed by end-points is a promising approach to improve efficiency, security, and robustness of communication networks in their various forms: To name a few examples, solutions based on end-point path selection have been proposed for routing on multiple optimality criteria [1], multi-tenant data centers [2], mobile ad-hoc networks [3], LEO satellite networks [4], intra-domain forwarding [5], and inter-domain forwarding [6]. However, proposals based on end-host path selection often encounter a stability concern: researchers have identified the problem that uncoordinated path-selection decisions by end-points may lead to persistent *oscillation*, i.e., an alternating grow-and-shrink pattern of traffic volumes on links [7,8]. The risk of oscillation still represents an obstacle to deployment of

* Corresponding author.

E-mail address: simon.scherrer@inf.ethz.ch (S. Scherrer).

path-aware networks [9] and gives rise to schemes that try to avoid oscillation [7,10–13]. While there is a rich literature presenting solutions for oscillation suppression, relatively little is known about *how exactly* and *by how much* instability from path selection deteriorates network performance. In other words, the solution to the oscillation problem is much clearer than both the impact vectors and the magnitude of the problem.

In this work, we therefore aim at qualifying and quantifying the effects of oscillatory path selection on various metrics of a network. To tackle this challenge, we must take into account that end-points in real path-aware networks employ algorithms which jointly perform path selection and congestion control (CC), i.e., multi-path congestion-control (MPCC) algorithms. In this work, we will focus on MPCC algorithms that are inspired by greedy, myopic path-selection behavior and thus simultaneously produce and react to oscillation. Furthermore, we require an analytical approach that (i) captures the congestion-window fluctuations that represent the oscillation, and (ii) is general enough to deliver fundamental insights into the nature of CC-assisted end-host path selection. Alas, fluid models [12,14–17] are well suited to represent equilibria in terms of the rough traffic distribution on a network; these models, however, either completely disregard congestion-control dynamics (such as the classic Wardrop model [14]) or fail to capture the small-scale dynamics of congestion-window fluctuations (as noted by Peng et al. who themselves operate with a fluid model [15]). More applied approaches, as employed in the design of multi-path TCP (MPTCP) [18–20], can capture oscillatory phenomena (e.g., the ‘flappiness’ of protocols [20]), but these approaches rely on ad-hoc reasoning from stylized network examples and experimental validation, which reduces their viability as generic analytic tools.

We argue that a so-called *axiomatic* approach recently initiated by Zarchy et al. [21] offers both the right analytical resolution and the required generality for the question at hand. This approach is axiomatic in a sense borrowed from economics and game theory, where properties with obvious desirability (e.g., the acyclicity of preferences [22] or the fairness of a bargaining outcome [23]) are formulated as axioms. Zarchy et al. apply this approach to congestion control by capturing desirable properties of CC protocols such as efficiency, fairness, and stability in axioms. The approach allows to analytically rate protocols with respect to these axioms and highlight the fundamental trade-offs between them. In our work, we further extend Zarchy et al.’s model to a multi-path context with the goal of characterizing fundamental properties of joint algorithms for path selection and CC.

1.1. Contribution

Our paper uses a theoretical model to investigate how network performance is affected by the instability due to greedy end-point path selection. In contrast to earlier theoretical models, we develop a model that is able to capture both path-selection dynamics and congestion-window fluctuations in Section 2. Within this model, we identify and formalize different classes of dynamic equilibria (in Sections 3 and 4) to which the flow dynamics can be expected to converge exponentially fast. These equilibria are essential for the analytical rating of MPCC protocols: In Section 5, we rate these dynamic equilibria with respect to a number of performance metrics (the axioms), which are inspired by the recently developed axiomatic approach to CC [21], but extended to accommodate path selection. This equilibrium formalization allows to derive the following insights in Section 6:

- **No trade-off between efficiency, convergence and loss avoidance:** Through appropriate protocol tuning, the metrics efficiency, loss avoidance, and convergence can be simultaneously optimized. Hence, there is no trade-off between these properties in theory.
- **Trade-off with fairness and responsiveness:** There is, however, a fundamental trade-off between the above metrics and the fairness and the responsiveness of a MPCC protocol. In particular, higher responsiveness makes a protocol less efficient, but more fair.
- **Effects of introducing end-host path selection:** By contrasting the axiomatic performance ratings for a general network with and without path selection, we obtain a multifaceted formalization of the performance impact of introducing end-host path selection. This formalization allows to interpret and quantify how unstable path selection affects network performance depending on network parameters. The insights gained from this approach show that there are both benefits and drawbacks of end-host path selection.

2. Model and assumptions


2.1. Discrete model

We leverage the analytical model of congestion control proposed by Zarchy et al. [21] and extend it to a multi-path context with path selection as illustrated in Fig. 1 (Notation summary in Table 1). In summary, N agents (denoted by set $A = [N] := \{0, \dots, N - 1\}$) compete for bandwidth on the bottleneck links of P parallel paths from set Π . Each agent $i \in A$ maintains a congestion window with size $cwnd_i$, which evolves over time t . At each moment $t \in \mathbb{N}_0$ in discrete time, any path $\pi \in \Pi$ accommodates a set $A_\pi(t)$ of agents that use path π at moment t , and carries load $f_\pi(t) = \sum_{i \in A_\pi(t)} cwnd_i(t)$. Moreover, in each time step t , every agent i takes two actions. First, agent i performs *congestion control*, i.e., adapts its congestion-window size $cwnd_i(t)$ according to a chosen CC protocol $CC_i(t, cwnd_i(t))$, resulting in congestion-window size $cwnd_i(t + 1)$. Second, agent i performs *path selection*, i.e., determines the path π such that $i \in A_\pi(t)$ and $cwnd_i(t)$ is included in $f_\pi(t)$, according to a given path-selection strategy. In Fig. 1 as well as in

Table 1

Notation used in our model in alphabetic order.

Symbol	Description
$A = [N]$	Set of agents in network
$A_\pi(t)$	Set of agents using path π at time t
$a_\pi(t)$	Number of agents using path π at time t
$\alpha(\tau)$	Additive increase given continuity time τ
β	Multiplicative-decrease parameter
C	Total bottleneck capacity of network
C_π	Bottleneck capacity of path π
$cwnd_i(t)$	Congestion-window size of agent i at time t
$f(t)$	Combined congestion-window size of all agents at time t
$f_\pi(t)$	Combined congestion-window size of all agents using path π at time t
$M_\pi(t)$	Set of agents who migrate away from path π at time t
m	Responsiveness (probability of switching to more attractive path in each time step)
N	Number of agents in the network
P	Number of paths in the network
Π	Set of paths in the network
$\pi_i(t)$	Path used by agent i at time t
$\pi_{\min}(t)$	Path with lowest utilization at time t
r	Reset softness (multiplicative decrease of congestion-window size on path switch)
$\text{rank}(\pi, t)$	Rank of path π at time t (number of paths with higher utilization than π at time t)
τ	Continuity time (time since last loss or path switch)
$\tau_i(t)$	Continuity time of agent i at time t
$z(a_\pi(t), N)$	Scaling factor for extrapolating on-migration flow volume from path flow

**Fig. 1.** Illustration of discrete model (Notation: $\bar{N} = N - 1$). The dotted arrow visualizes path migration by agent 1 from path π to path π'' .

our following analysis, agents implement probabilistic greedy path selection, i.e., switch to the path carrying the lowest load in the last time step with a given probability m . Finally, in order to investigate different behaviors for congestion-window adaptation upon path switches, we introduce a reset-softness parameter $r \in [0, 1]$ that determines the extent of congestion-window reduction for path-switching agents (e.g., agent 1 in Fig. 1).

The agents are further constrained by path capacities C_π , $\pi \in \Pi$, where C_π is the amount of data in maximum segment size (MSS) that can be transmitted on path π during one round-trip time (RTT). If the capacity C_π of path π is exceeded by the flow $f_\pi(t)$, the agents $A_\pi(t)$ experience packet loss and take this loss into account in their congestion-control protocol.¹ For example, the TCP Reno protocol, with a multiplicative decrease of 0.5 as a reaction to loss and an additive increase of 1 otherwise, is modeled as follows for an agent i using path π at time t :

$$\text{TCPReno}(t, cwnd_i(t)) = \begin{cases} cwnd_i(t) + 1 & \text{if } f_\pi(t) \leq C_\pi \\ 0.5 \cdot cwnd_i(t) & \text{otherwise} \end{cases} \quad (1)$$

¹ We note that this loss modeling is a simplification in three respects. First, loss may already occur when $f_\pi(t) > s_\pi$, namely if all agents send out all traffic $f(t)$ in a burst that exceeds the buffer size s_π . Second, even if $f_\pi(t) > C_\pi$, the loss may not be perceived by all agents. Third, CC algorithms may react differently depending on the number of recent losses.

2.2. Scenario of interest and assumptions

Since the goal of this work is to characterize the worst-case effects of oscillatory path selection, our analysis throughout the paper will focus on a network scenario that maximizes the severity of load oscillation. This scenario has the following properties, which henceforth serve as assumptions:

Greedy load-adaptive path selection. Oscillation is caused by greedy, myopic path selection behavior [13], which dynamically determines the number $a_\pi(t) = |A_\pi(t)|$ of agents on path π . In any time step t , agents seek out the path $\pi_{\min}(t)$ with the lowest bottleneck utilization $f_\pi(t)/C_\pi$ and hence the lowest latency (assuming roughly equal propagation delay of all paths as stated below) and lowest loss rate. Since monitoring the state of alternative paths and switching paths consume resources, agents may not consider a path change in every time step. Instead, the path-selection behavior is regulated by a *path-migration probability* $m \in (0, 1]$, denoting the probability with which an agent switches to a more attractive path in any time step. Alternatively, m can be interpreted as a measure for the *responsiveness* of agents.

Sequential multi-path usage. The intensity of oscillations grows with the size of shifted flow volume per time unit. In order to maximize oscillation, we therefore assume that a path-switching agent completely stops using its previously used path and exclusively sends on the newly selected path. This coarse-granular migration behavior produces sequential instead of concurrent usage of multiple paths. This mode of sequential multi-path usage approximates the actual behavior of real-world algorithms such as MPTCP, which tends to use only the most attractive path for data transmission and sends a negligible amount of probing traffic over the alternative paths [12,18,20]. Moreover, the average utility improvement per user that is possible by concurrently using multiple paths instead of a single selected path vanishes for a high number of agents [24]. Sequential multi-path usage implies that $\sum_{\pi \in \Pi} a_\pi(t) = N \forall t$.

Disjoint and similar paths. We investigate a network consisting of paths that are parallel, disjoint and equal in terms of latency D_π and bottleneck capacity $C_\pi = C/P$, where C is the bottleneck capacity of the complete network. Such a network, while being a simplification of general networks, is likely to bring out the worst-case effects of myopic, greedy path selection, which are the subject of this paper. In particular, load oscillations are strongest if the actions of the sending agents are strongly correlated because they react to the same (potentially misleading) feedback signals (i.e., path loss and latency) simultaneously [13]. If agents sharing a link react to different feedback signals or at different times, e.g., because they are using different paths with different round-trip latencies, their actions are less strongly correlated and the flow dynamics are likely to oscillate less. The feedback synchronization by equal path RTTs also ensures that the discrete time steps of the model have consistent duration across all paths.

2.3. Stochastic dynamics

In summary, a multi-path congestion-control protocol $MPCC(CC, m, r)$ is a combination of a CC protocol $CC(t)$, a responsiveness parameter m , and a reset-softness parameter r . In a network with path selection, the MPCC dynamics can thus be represented by a pair of functions $(a_\pi(t), f_\pi(t))$ for any path $\pi \in \Pi$. Since the path-selection behavior is probabilistic (regulated by responsiveness parameter m), the MPCC dynamics are not uniquely determined by initial conditions, but need to be modeled as a stochastic process. In particular, the MPCC dynamics under universal adoption of $MPCC(CC, m, r)$, are given by

$$a_\pi(t+1) = \begin{cases} a_\pi(t) - |M_\pi(t)| & \text{if } \pi \neq \pi_{\min}(t) \\ a_\pi(t) + \sum_{\tilde{\pi} \neq \pi} |M_{\tilde{\pi}}(t)| & \text{otherwise} \end{cases} \quad (2a)$$

$$f_\pi(t+1) = \begin{cases} f_\pi(t) - \sum_{i \in M_\pi(t)} cwnd_i(t) + \sum_{i \in A_\pi(t) \setminus M_\pi(t)} \Delta cwnd_i(t) & \text{if } \pi \neq \pi_{\min}(t) \\ f_\pi(t) + \sum_{\tilde{\pi} \neq \pi} \sum_{j \in M_{\tilde{\pi}}(t)} r \cdot cwnd_j(t) + \sum_{i \in A_\pi(t)} \Delta cwnd_i(t) & \text{otherwise,} \end{cases} \quad (2b)$$

where $M_\pi(t)$ is a random subset of $A_\pi(t)$, which contains the agents who leave path π at time t , and $\Delta cwnd_i(t) = cwnd_i(t+1) - cwnd_i(t)$. Intuitively, the flow on a more congested path π is reduced by the congestion windows of all agents M_π that leave the path, and increased by the congestion-window growth of the remaining agents $A_\pi(t) \setminus M_\pi(t)$. In contrast, the flow on the least congested path π is increased by the reset congestion-window sizes $r \cdot cwnd_j(t)$ of the agents $j \in M_{\tilde{\pi}}$ who migrate to path π and the congestion-window growth of the previously present agents $A_\pi(t)$.

2.4. Expected dynamics

While the formulations in Eq. (2) capture the evolutionary dynamics of an MPCC system, their discrete and probabilistic nature hinders analytic treatment. However, as we investigate large-scale systems with a high number of agents, the law of large numbers allows that the probabilistic elements in Eq. (2) can be well approximated by their expected values and traffic randomness can be greatly ignored. For the remainder of this paper, we therefore consider the *expected* MPCC dynamics, where the recursion on the random variables $(a_\pi(t), f_\pi(t))$ is approximated with a recursion on the expectations $(\hat{a}_\pi(t), \hat{f}_\pi(t))$ (where we write $\hat{x} := \mathbb{E}[x]$ for any function x). The accuracy of this approximation will be validated with simulations in Appendix B.

Concerning the agent dynamics in Eq. (2a), we note that $\mathbb{E}[|M_\pi(t)|] = m \cdot \hat{a}_\pi(t)$ for any path $\pi \neq \pi_{\min}(t)$. Moreover, the expected volume of flow associated with the agents in M_π in Eq. (2b) is a proportional share of the expected total flow $\hat{f}_\pi(t)$

on path π : $\mathbb{E}[\sum_{i \in M_\pi(t)} \text{cwnd}_i(t)] = m \cdot \hat{f}_\pi(t)$. By the same argument, it holds that $\mathbb{E}[\sum_{\tilde{\pi} \neq \pi} \sum_{j \in M_{\tilde{\pi}}(t)} r \cdot \text{cwnd}_j(t)] = m \cdot r \cdot \sum_{\tilde{\pi} \neq \pi} \hat{f}_{\tilde{\pi}}(t)$ for $\pi = \pi_{\min}(t)$. However, in order to make the second case of Eq. (2b) independent of flows $f_{\tilde{\pi}}$ on alternative paths, we additionally make the following approximation: $\sum_{\tilde{\pi} \neq \pi} \hat{f}_{\tilde{\pi}}(t) \approx (N - \hat{a}_\pi(t)) / \hat{a}_\pi(t) \cdot \hat{f}_\pi(t) = z(\hat{a}_\pi(t), N) \cdot \hat{f}_\pi(t)$, where $z(\hat{a}_\pi(t), N)$ is henceforth referred to as the *extrapolation factor*. In this approximation, the flow on path π is scaled proportionally to the number of agents $N - \hat{a}_\pi(t)$ on other paths. This approximation can be justified on the grounds that in a steady state, imbalances in path load are likely to stem from imbalances in the number of agents between paths, not from imbalances in the average congestion-window size between paths.

Finally, in order to arrive at the expected flow dynamics $\hat{f}_\pi(t)$, the expected combined congestion-window change $\mathbb{E}[\sum_{k \in A_\pi(t)} \Delta \text{cwnd}_k(t)]$ (or for $A_\pi(t+1)$, respectively) must be formalized. Of course, this change depends on the CC protocols employed by the agents. In order to maximize the generality of our analysis, we rely on the following generic form of a loss-based CC protocol employed by each agent i , where $\pi_i(t)$ denotes the path that agent i uses at time t :

$$\text{CC}_i(t, \text{cwnd}_i(t)) = \begin{cases} \text{cwnd}_i(t) + \alpha(\tau_i(t)) & \text{if } f_{\pi_i(t)}(t) \leq C_{\pi_i(t)} \\ \beta \cdot \text{cwnd}_i(t) & \text{otherwise} \end{cases} \quad (3)$$

Here, $\tau_i(t)$ is the so-called *continuity time* of agent i , i.e., the number of time steps in which agent i has already been on its current path without experiencing packet loss. This continuity time is the argument to a function α , which determines the additive increase to the congestion window in absence of loss. This formulation allows to mimic the window-growth behavior in classic TCP Reno [25], in the widely deployed TCP CUBIC [26], in the slow-start phase of many TCP protocols [27], or in more theoretical MIMD protocols [28]. Finally, $\beta \in [0, 1]$ is a parameter that determines the multiplicative decrease of the congestion-window size in the case of packet loss, which is the predominant practice in CC protocols.

Based on the probability distribution for the continuity time $\tau_i(t)$ of any agent $i \in A_\pi(t)$ at time t from Appendix A, we can calculate the average congestion-window increase per agent conditioned on the path π used by the agent at time t : $\hat{\alpha}_\pi(t) = \sum_{\tau=0}^{\infty} \mathbb{P}[\tau_i(t) = \tau \mid i \in A_\pi(t)] \cdot \alpha(\tau)$. This average congestion-window increase then allows to obtain the aggregate additive increase in absence of loss. In contrast, loss reduces the expected flow volume $\hat{f}_\pi(t)$ through multiplicative decrease β , complementing the effects of out-migration (for $\pi \neq \pi_{\min}(t)$) or in-migration (for $\pi_{\min}(t)$). Under universal adoption of a protocol $\text{MPCC}(\text{CC}, m, r)$, the expected dynamics therefore are:

$$\hat{a}_\pi(t+1) = \begin{cases} (1-m) \cdot \hat{a}_\pi(t) & \text{if } \pi \neq \pi_{\min}(t) \\ (1-m) \cdot \hat{a}_\pi(t) + m \cdot N & \text{otherwise} \end{cases} \quad (4a)$$

$$\hat{f}_\pi(t+1) = \begin{cases} (1-m) \cdot \hat{f}_\pi(t) + \hat{\alpha}_\pi(t) \cdot (1-m) \cdot \hat{a}_\pi(t) & \text{if } \pi \neq \pi_{\min}(t) \wedge \hat{f}_\pi(t) \leq C_\pi \\ (1+m \cdot r \cdot z(\hat{a}_\pi(t), N)) \cdot \hat{f}_\pi(t) + \hat{\alpha}_\pi(t) \cdot \hat{a}_\pi(t) & \text{if } \pi = \pi_{\min}(t) \wedge \hat{f}_\pi(t) \leq C_\pi \\ \beta \cdot (1-m) \cdot \hat{f}_\pi(t) & \text{if } \pi \neq \pi_{\min}(t) \wedge \hat{f}_\pi(t) > C_\pi \\ (\beta + m \cdot r \cdot z(\hat{a}_\pi(t), N)) \cdot \hat{f}_\pi(t) & \text{if } \pi = \pi_{\min}(t) \wedge \hat{f}_\pi(t) > C_\pi \end{cases} \quad (4b)$$

2.5. Limitations

While our model presents a tractable approach to analyze oscillatory MPCC dynamics, our investigation and the resulting insights have clear limitations worth addressing in future research. In particular, as our network model is an extension of the network model by Zarchy et al. [21], our work inherits some limitations noted by Zarchy et al. most importantly the assumption of synchronized feedback, the focus on a specific type of network, and the disregard for queuing dynamics. However, it is noteworthy that our work addressed the previously identified challenge concerning randomized protocols through the concept of expected dynamics. In general, the comprehensiveness of our analysis would benefit from relaxing the worst-case conditions elicited in Section 2.2, most prominently the assumption of disjoint and similar paths, and from introducing latency-based and model-based CC protocols.

3. Lossless equilibria

In order to rate MPCC protocols, we focus on the *equilibria* that these protocols induce, i.e., stable load patterns to which the MPCC dynamics from Eq. (4) eventually converge. In this section, we characterize one class of equilibria that are attained before the capacity limit of any bottleneck link is exceeded, i.e., these equilibria are *lossless*. Equilibria without this lossless property, i.e., *lossy equilibria*, are presented in Section 4. All of these equilibria are *dynamic equilibria*, i.e., periodic patterns of the number of agents and the load on the different paths. Note that the insights regarding equilibria only apply to the theoretical construct of expected dynamics in an exact sense, and only approximately apply to actual MPCC dynamics.

3.1. Structure of lossless equilibria

In order to characterize lossless equilibria, we need to investigate whether the expected MPCC dynamics tend to exhibit a certain pattern in the case where capacity limits are disregarded. Unfortunately, even this simplified discrete dynamical system (determined by Eq. (4) without the two last cases of Eq. (4b)) is analytically intractable due to the presence of case distinctions in the evolution functions [29]. Instead, we use a hybrid approach, similar to previous work [30]: By performing simulations as in Fig. B.10, we arrive at the following two observations about MPCC dynamics with greedy, myopic agents sharing parallel and similar paths (cf. Section 2.2), which serve as a basis for further analytical investigation:

In-migration is utilization-maximizing: Whenever path π with minimal utilization within the expected dynamics, i.e., $\hat{u}_\pi(t) = \hat{f}_\pi(t)/(C/P)$, experiences in-migration according to the second case of Eq. (4b), this path tends to become the most utilized path in the next time step.²

Out-migration is order-preserving: If two paths π and $\tilde{\pi}$ with $\hat{u}_\pi(t) > \hat{u}_{\tilde{\pi}}(t)$ experience out-migration according to the first case of Eq. (4b), it tends to hold that $\hat{u}_\pi(t+1) > \hat{u}_{\tilde{\pi}}(t+1)$.

If the expected dynamics consistently conform to these two observations, they exhibit the following pattern which uniquely determines the least utilized path in every time step:

Definition 1. MPCC dynamics exhibit ***P-step oscillation*** if there exists a time $t_0 \geq 0$ such that

$$\forall T \geq 0. \quad \text{rank}(\pi, t_0) = p \implies \text{rank}(\pi, t_0 + T) = (p + T) \bmod P, \quad (5)$$

where $\text{rank}(\pi, t)$ ranks all paths $\pi \in \Pi$ in descending order according to their utilization at time t :

$$\text{rank}(\pi, t) = p \iff |\{\tilde{\pi} \mid \hat{u}_{\tilde{\pi}}(t) > \hat{u}_\pi(t)\}| = p. \quad (6)$$

In P -step oscillation, the assignment of the rank to paths changes in a round-robin fashion, i.e., in any time step t , every path π rises by one rank, except the path with rank $P - 1$ (i.e., with the lowest expected utilization), which obtains rank 0 at time $t + 1$. After P time steps, a path reaches its original place in the ranking order, i.e., $\text{rank}(\pi, t) = \text{rank}(\pi, t + P)$ for all $t \geq t_0$. We present an argument for the prevalence of P -step oscillation in Section 3.3.

3.2. Lossless agent equilibrium

As this P -step oscillation uniquely determines the least congested path in any time step $t \geq t_0$, this pattern also determines the agent-migration dynamics. Starting from an agent distribution $\{a_\pi(t_0)\}_{\pi \in \Pi}$ at time t_0 , all the paths π with $\text{rank}(\pi, t_0) \neq P - 1$ will experience an outflow of agents (according to case 1 in Eq. (4a)) and only the path with rank $P - 1$ experiences an inflow of agents (according to case 2 in Eq. (4a)). In a single round of P -step oscillation with start time t_0 , the path $\pi^{(0)}$ with $\text{rank}(\pi^{(0)}, t_0) = 0$ will thus first experience agent outflow for $P - 1$ times and then once experience agent inflow. Hence, the following difference equation characterizes the discrete dynamical system for a granularity of P time steps:

$$\hat{a}_{\pi^{(0)}}(t_0 + P) = (1 - m)^P \cdot \hat{a}_{\pi^{(0)}}(t_0) + m \cdot N, \quad (7)$$

To find an equilibrium of the dynamic system for the agent dynamics on $\pi^{(0)}$, we identify a fixed point of the difference equation in Eq. (7), i.e., we solve

$$\hat{a}^{(0)} = (1 - m)^P \cdot \hat{a}^{(0)} + m \cdot N \iff \hat{a}^{(0)} = \frac{m \cdot N}{1 - (1 - m)^P}, \quad (8)$$

where $\hat{a}^{(0)}$ is the *equilibrium value* for any $\hat{a}_\pi(t)$ with $\text{rank}(\pi, t_0) = 0$, which generalizes as follows:

Insight 1 (Convergence to Unique Dynamic Agent Equilibrium). Under P -step oscillation, the expected agent dynamics $\{\hat{a}_\pi(t)\}_{\pi \in \Pi}$ of an MPCC system asymptotically converge to a unique dynamic equilibrium, i.e., a cyclic series of states. This dynamic equilibrium of the agent dynamics consists of P states in each of which the rank- p path accommodates the corresponding equilibrium amount of agents $\hat{a}^{(p)}$, i.e.,

$$\hat{a}_\pi(t) = \hat{a}^{(\text{rank}(\pi, t))}, \text{ where } \hat{a}^{(p)} = \frac{(1 - m)^P \cdot m \cdot N}{1 - (1 - m)^P}. \quad (9)$$

This convergence can be shown by finding a *trajectory function*:

² This observation suggests that myopic, greedy load-adaptive path selection is not a Nash equilibrium strategy, which has also been demonstrated by recent research [13].




Fig. 2. Convergence to lossless agent equilibrium for $P = 3$ and $m = 0.1$.




Fig. 3. Convergence to lossless flow equilibrium for $P = 3$, $m = 0.1$, and $r = 0.5$.

Definition 2. A **trajectory function** $x_\pi^{(p)}(t)$ is an explicit interpolation function that yields the correct value of path-specific dynamics $x_\pi(t)$ at all moments where path π has rank p :

$$\forall k \in \mathbb{N}_{\geq 0}. \quad x_\pi^{(p)}(t_{\pi p} + k \cdot P) = x_\pi(t_{\pi p} + k \cdot P), \tag{10}$$

where $t_{\pi p} = \min\{t \mid t \geq t_0 \wedge \text{rank}(\pi, t) = p\}$ and t_0 is the start time of P -step oscillation.

For the agent dynamics $\hat{a}_\pi(t)$, such a trajectory function is given by

$$\hat{a}_\pi^{(p)}(t) = (\hat{a}_\pi(t_{\pi p}) - \hat{a}^{(p)}) \cdot (1 - m)^{t - t_{\pi p}} + \hat{a}^{(p)}. \tag{11}$$

As $\lim_{t \rightarrow \infty} \hat{a}_\pi^{(p)}(t) = \hat{a}^{(p)}$, the trajectory functions converge to the equilibrium found above exponentially fast. Fig. 2 visualizes the asymptotic convergence to the dynamic equilibrium $\{\hat{a}^{(p)}\}_{p \in [P]}$ (highlighted in blue) along the trajectory functions.

3.3. Lossless flow equilibrium

After identifying the agent equilibrium in Section 3.2, we identify the equilibria of the MPCC flow dynamics $\{\hat{f}_\pi(t)\}_{\pi \in \Pi}$ in this section. We first consider *hypothetical equilibria*, which are equilibria of the flow dynamics under the assumption that the capacity of each path is never exceeded. In a second step, we will show under which conditions these hypothetical equilibria are actual equilibria.

3.3.1. Hypothetical flow equilibria

To find the hypothetical equilibria of the flow dynamics, we can simplify the flow dynamics from Eq. (4b) by disregarding the capacity limit C/P . In addition, we insert the equilibrium agent levels $\hat{a}^{(p)}$ from Section 3.2 and the expected additive increase $\hat{\alpha}^{(p)}$ derived in Appendix A to arrive at the following formulation:

$$\hat{f}_\pi(t+1) = \begin{cases} (1-m) \cdot (\hat{f}_\pi(t) + \hat{\alpha}^{(\text{rank}(\pi,t))} \cdot \hat{a}^{(\text{rank}(\pi,t))}) & \text{if } \text{rank}(\pi, t) \neq P-1 \\ (1+m \cdot r \cdot z(m, P)) \cdot \hat{f}_\pi(t) + \hat{\alpha}^{(P-1)} \cdot \hat{a}^{(P-1)} & \text{if } \text{rank}(\pi, t) = P-1, \end{cases} \quad (12)$$

where the extrapolation factor z is only dependent on m and P given the agent equilibrium, i.e., $z(m, P) = N/\hat{a}^{(P-1)} - 1 = (1 - (1-m)^{P-1})/(m \cdot (1-m)^{P-1})$.

Similar to Eq. (7), we set up a first-order difference equation for the dynamics for the path that has rank p at time t_0 (where the P -step oscillation starts) and find a fixed point that is attained every P time steps, for example for ranks 0 and $P-1$:

$$\hat{f}^{(0)} = \frac{((1+m \cdot r \cdot z(m, P)) \cdot (\sum_{p=0}^{P-2} \hat{\alpha}^{(p)}) + \hat{\alpha}^{(P-1)}) \cdot \hat{a}^{(P-1)}}{1 - (1+m \cdot r \cdot z(m, P)) \cdot (1-m)^{P-1}}, \quad (13a)$$

$$\hat{f}^{(P-1)} = \frac{(\sum_{p=0}^{P-2} \hat{\alpha}^{(p)} + \hat{\alpha}^{(P-1)}) \cdot (1-m)^{P-1} \cdot \hat{a}^{(P-1)}}{1 - (1+m \cdot r \cdot z(m, P)) \cdot (1-m)^{P-1}}. \quad (13b)$$

The fixed point for a general rank p can be derived analogously and expressed by a similar (albeit quite complicated) term $\hat{f}^{(p)}$ shown in Eq. (C.1) in Appendix C. These fixed points $\{\hat{f}^{(p)}\}_{p \in [P]}$ constitute the hypothetical equilibrium, i.e., if a rank- p path carries flow volume $\hat{f}^{(p)}$, the path will carry this flow volume again P time steps later, where it is again the rank- p path.

Insight 2 (Hypothetical Dynamic Flow Equilibrium). *If capacity limits of links are disregarded, the dynamic equilibrium of the flow dynamics $\{\hat{f}_\pi(t)\}_{\pi \in \Pi}$ consists of P states in each of which the rank- p path accommodates flow volume $\hat{f}^{(p)}$.*

In order for such an equilibrium to be valid, it must be consistent with P -step oscillation, i.e., it must hold that $\hat{f}^{(p)} > \hat{f}^{(p+1)}$ for all $p \in [P-1]$. Interestingly, if a certain parameter combination is associated with an invalid equilibrium, it follows that P -step oscillation is fundamentally impossible for that parameter combination. However, we show in Appendix C that only a small part of the parameter space, containing rather extreme parameters, is inconsistent with P -step oscillation.

Similarly as in Section 3.2, convergence to this equilibrium can be proven using a trajectory function (cf. Definition 2). The following trajectory function yields the correct flow volume in all subsequent time steps where path π has rank p again:

$$\hat{f}_\pi^{(p)}(t) = (\hat{f}_\pi(t_{\pi p}) - \hat{f}^{(p)}) \cdot ((1+m \cdot r \cdot z(m, P)) \cdot (1-m)^{P-1})^{\frac{t-t_{\pi p}}{P}} + \hat{f}^{(p)}. \quad (14)$$

The limit of this trajectory function for $t \rightarrow \infty$ is the equilibrium value $\hat{f}^{(p)}$, which establishes convergence;³ this is illustrated in Fig. 3.

3.3.2. Actual flow equilibrium

Intuitively, this hypothetical equilibrium given by $\{\hat{f}^{(p)}\}_{p \in [P]}$ is an actual equilibrium of the MPCC dynamics if the convergence is not disturbed by the capacity limit C/P on any path π , i.e., if the trajectory functions for all ranks consistently remain below C/P . We therefore require an upper bound on all trajectory functions $\{\hat{f}_\pi^{(p)}(t)\}_{p \in [P]}$. Thanks to the structure of P -step oscillation, it holds that $\hat{f}^{(p)} > \hat{f}^{(p+1)} \forall p \in [P-1]$. Therefore, in the hypothetical equilibrium, $\hat{f}^{(0)}$ represents an upper bound on the flow dynamics. We speak of flow dynamics with *consistent trajectories* if such an ordering not only holds on the equilibrium values $\hat{f}^{(p)}$, but also on the trajectory functions $\hat{f}_\pi^{(p)}(t)$ for all paths π :

Definition 3. Flow dynamics $\{\hat{f}_\pi(t)\}_{t \geq 0}$ have **consistent trajectories** at time point t' if on every path $\pi \in \Pi$, the rank-specific trajectory functions $\{\hat{f}_\pi^{(p)}(t)\}_{p \in [P]}$ satisfy the following condition:

$$\forall p \in [P-1], t \geq t'. \quad \hat{f}_\pi^{(p)}(t) > \hat{f}_\pi^{(p+1)}(t) \quad (16)$$

³ Note that $\hat{f}^{(0)}$ from Eq. (13a) is undefined for $r = 1$, as the flow dynamics do not converge to a fixed point in that case. Given $r = 1$, the trajectory function for rank 0 can be expressed with the following linear function, which has no limit:

$$\hat{f}_\pi^{(0)}(t) = \left[(1-m)^{1-P} \cdot \left(\sum_{p=0}^{P-2} \hat{\alpha}^{(p)} \right) + \hat{\alpha}^{(P-1)} \right] \cdot \hat{a}^{(P-1)} \cdot P^{-1} \cdot (t - t_{\pi 0}) + \hat{f}_\pi(t_{\pi 0}). \quad (15)$$




Fig. 4. Type-1 lossy equilibrium for $P = 3$, $m = 0.45$, and $r = 0.9$ (One period is highlighted in light-blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

As trajectories are always eventually consistent, the trajectory function $\hat{f}_\pi^{(0)}(t)$ for rank 0 is therefore an upper bound on all trajectory functions $\{\hat{f}_\pi^{(p)}(t)\}_{p \in [P]}$ and by consequence also an upper bound on the flow dynamics $\{\hat{f}_\pi(t)\}_{t \geq 0}$ for any path π . As $\hat{f}_\pi^{(0)}(t)$ is monotonic, its function values will not exceed C/P if $\hat{f}_\pi(t_{\pi p}) \leq C/P$ and $\hat{f}_\pi^{(0)} \leq C/P$. Due to the introduction of capacity limits, it is necessary to alter the definition of $t_{\pi p}$ to be the first point in time after oscillation began (at t_0) where $\text{rank}(\pi, t_{\pi p}) = p$ (as before) and additionally $\hat{f}_\pi(t_{\pi p}) \leq C/P$.⁴ Therefore, we arrive at the following insight:

Insight 3 (Dynamic Lossless Flow Equilibrium). *The hypothetical equilibrium (disregarding capacity limitations) from Insight 2 is an actual, lossless equilibrium (taking capacity limits into account) for the flow dynamics $\{\hat{f}_\pi(t)\}_{\pi \in \Pi}$ if and only if $\hat{f}_\pi^{(0)} \leq C/P$, i.e., the maximum flow-equilibrium level does not exceed the bottleneck capacity of any path.*

4. Lossy equilibria

In this section, we characterize *lossy* equilibria, i.e., dynamic equilibria where $\hat{f}_\pi^{(0)} > C/P$ and the flow dynamics therefore periodically exceed bottleneck capacities.

4.1. Structure of lossy equilibria

In order to identify the typical structure of lossy equilibria, we again rely on simulations similar to Section 3.1. Based on these simulations, we can distinguish two types of lossy equilibria, illustrated in Figs. 4 and 5. Note that both of these lossy-equilibrium types are characterized by flow volumes $\{\hat{f}_\pi^{(p)*}\}_{p \in [P]}$, each carried by the path with rank p in the state that is designated as the initial state of the lossy equilibrium ($t = 0$ in the figures) and is periodically revisited every L time steps. Moreover, the boundary points, i.e., the largest and smallest flow volume arising in a lossy equilibrium, are denoted by \hat{f}^{*+} and \hat{f}^{*-} , respectively.

The main distinguishing property of type-1 lossy equilibria (cf. Fig. 4) is that these lossy equilibria are consistent with P -step oscillation despite the occasional multiplicative decrease β on rank-0 paths. In contrast, type-2 lossy equilibria (cf. Fig. 5) temporarily deviate from P -step oscillation whenever there is packet loss on a path. In that case, the rank-0 path with loss directly becomes the rank- $(P - 1)$ path in the subsequent time step. However, even in type-2 lossy equilibria, P -step oscillation eventually resumes, e.g., at $t = 2$ in Fig. 5. Type-1 equilibria typically appear for a relatively high migration rate m , whereas type-2 equilibria tend to appear for lower migration rates.

4.2. Flow equilibria

While the lossy equilibria cannot be characterized as simply as the lossless equilibria in Section 3.3, it is feasible to determine the flow-volume bounds for the presented types of lossy equilibria.

⁴ We assume that such a $t_{\pi p}$ always exists as any reasonable CC's reaction to loss reduces $\hat{f}_\pi(t)$ below C/P eventually.




Fig. 5. Type-2 lossy equilibrium for $P = 3$, $m = 0.1$, and $r = 1$ (One period is highlighted in light-blue).

Regarding the upper bound, the central question is how high the upper boundary point \hat{f}^{*+} can become. Both for type-1 and type-2 lossy equilibria, we note that \hat{f}^{*+} is reached after one round of P -step oscillation starting from a flow volume below the capacity limit. Hence, an upper bound on \hat{f}^{*+} can be represented as follows:

$$\hat{f}^{*+} \leq \hat{f}_{\pi}^{(0)}(t_{\pi 0} + P) \text{ where } \hat{f}_{\pi}^{(p)}(t_{\pi 0}) = C/P. \quad (17)$$

The trajectory function from Eq. (14) (or from Eq. (15) for $r = 1$) is used to calculate the effects of one round of P -step oscillation on flow volume C/P , which is the highest flow volume from which ordinary P -step oscillation can proceed. Note that this trajectory function is only usable if the agent dynamics are in equilibrium according to Section 3.2. Type-1 lossy equilibria preserve P -step oscillation and thus also the corresponding agent equilibria. For type-2 lossy equilibria, however, P -step oscillation is occasionally disturbed, which can result in agent dynamics out of equilibrium. However, multiple rounds of P -step oscillation precede the moment of reaching \hat{f}^{*+} and the convergence to the agent equilibrium is exponential. Hence, we observe that the agent dynamics are close to the agent equilibrium and the trajectory function can therefore be used to obtain an approximate upper bound for type-2 lossy equilibria.

Regarding the lower bound, we now investigate how low the lower boundary point \hat{f}^{*-} can become. In type-1 lossy equilibria, \hat{f}^{*-} is reached after $P - 1$ time steps with agent outflow on an overloaded rank-0 path. Combined with the multiplicative decrease β in the first of these $P - 1$ time steps, we can thus formulate a lower bound on \hat{f}^{*-} for type-1 lossy equilibria:

$$\hat{f}^{*-} > \beta \cdot (1 - m)^{P-1} \cdot C_{\pi} \quad (18)$$

given a rank-0 path that is only infinitesimally overloaded and $\alpha(\tau) > 0, \forall \tau \in \mathbb{N}_{>0}$. For type-2 lossy equilibria, this lower bound is too pessimistic, as the combination of multiplicative decrease and agent out-migration directly transforms the overloaded rank-0 path into the least utilized path and there are no further consecutive time steps with agent out-migration on this path. Hence, $\beta \cdot (1 - m) \cdot C_{\pi}$ suffices as a lower bound for type-2 lossy equilibria. For a validation of these lower bounds by simulations, consult Fig. D.14 in Appendix D.

5. Axioms

In this section, we use an axiomatic approach inspired by Zarchy et al. [21] to derive insights regarding the effects of oscillatory path selection. We adapt a number of their axioms, which were formulated for a single-path context, to a multi-path context in Section 5.1. In Section 5.2, we evaluate the equilibria from Sections 3 and 4 with respect to these axioms.

5.1. List of axioms

In our axiomatic approach to multi-path congestion control, axioms correspond to desirable properties that MPCC protocols should possess. However, as these properties refer to general and vague concepts (e.g., efficiency or fairness), the conditions for possessing these properties are usually not well-defined. Therefore, the axioms here are formalized as metrics for rating an MPCC protocol with respect to a certain property, instead of binary indicators of whether the protocol possesses the given property. Concretely, we consider the following axioms in this work:

Axiom 1 (Efficiency). An MPCC protocol is ϵ -efficient if under universal adoption of this protocol, the bottleneck utilization of every path π with capacity C/P is never lower than a share ϵ after some time t' :

$$\exists t'. \quad \forall t \geq t', \pi \in \Pi. \quad \frac{P \cdot \hat{f}_\pi(t)}{C} \geq \epsilon \quad (19)$$

Larger values of ϵ are better, and we consider an ϵ -efficient protocol optimal if $\epsilon \geq (C - s)/C$, where s is the buffer size.⁵

Axiom 2 (Loss Avoidance). An MPCC protocol is λ -loss-avoiding if under universal adoption, the loss rate on any path π with capacity C/P never exceeds λ after some time t' :

$$\exists t'. \quad \forall t > t', \pi \in \Pi. \quad \frac{\hat{f}_\pi(t)}{C/P} - 1 \leq \lambda \quad (20)$$

Thus, smaller values of λ are better, and a 0-loss-avoiding protocol is optimal.

Axiom 3 (Convergence). An MPCC protocol is γ -convergent if under universal adoption, the flow volume $\hat{f}_\pi(t)$ on every path π lies consistently within a range $[\gamma \cdot \hat{f}_\pi^+, \hat{f}_\pi^+]$ below a path-specific maximum level \hat{f}_π^+ after some time t' :

$$\exists t' > 0, \hat{f}_\pi^+. \quad \forall t > t', \pi \in \Pi. \quad \gamma \cdot \hat{f}_\pi^+ \leq \hat{f}_\pi(t) \leq \hat{f}_\pi^+ \quad (21)$$

Thus, larger values of γ are better, and a 1-convergent protocol is optimal.

Axiom 4 (Fairness). An MPCC protocol is η -fair if under universal adoption, the variance of congestion-window sizes of all agents $i \in A$ in the network never exceeds η after some time t' :⁶

$$\exists t' > 0. \quad \forall t > t'. \quad \text{Var}_{i \in A}[cwnd_i(t)] \leq \eta \quad (22)$$

Thus, smaller values of η are better, and a 0-fair protocol is optimal.

For any axiom metric μ , we write $\mu(\text{MPCC})$ for the most desirable value of metric μ that the protocol MPCC can be rated with.

5.2. Axiomatic characterization of equilibria

The axioms in Section 5.1 refer to characteristics which are eventually attained and then persistently preserved by the flow dynamics. Hence, a natural way to axiomatically rate an MPCC protocol is to evaluate the equilibria (i.e., stable states) of this protocol (cf. Sections 3 and 4).

Efficiency (Axiom 1). We distinguish lossless and lossy flow equilibria. If there is a lossless equilibrium ($\hat{f}^{(0)} \leq C/P$), the minimal flow volume ever carried by any path π is the equilibrium value for rank $P - 1$, i.e., $\hat{f}^{(P-1)}$.⁷ The network-wide efficiency level is therefore $\epsilon = P \cdot \hat{f}^{(P-1)}/C$. In contrast, for lossy equilibria, the efficiency level is the lower bound on the lower boundary point \hat{f}^{*-} according to Section 4.2. Depending on the lossy-equilibrium type, this lower bound is given by $\beta \cdot (1 - m)^{P-1}$ or $\beta \cdot (1 - m)$, respectively. Since the lower bound is never higher for type 1 than for type 2, we consider $\beta \cdot (1 - m)^{P-1}$ to be the minimum flow volume for lossy equilibria. While this lower bound is too pessimistic for lossy equilibria of type 2, these type-2 lossy equilibria mostly appear for low values of m , where the difference between the two bounds is small.

$$\epsilon(\text{MPCC}(\alpha, \beta, m, r)) \geq \begin{cases} P \cdot \hat{f}^{(P-1)}/C & \text{if } \hat{f}^{(0)} \leq C/P \\ \beta \cdot (1 - m)^{P-1} & \text{otherwise} \end{cases} \quad (23)$$

Loss avoidance (Axiom 2). If all paths are in lossless equilibrium ($\hat{f}^{(0)} \leq C/P$), it is clear that the maximum loss rate in the whole network is 0. If the network is in lossy equilibrium, the maximum loss rate is determined by the upper boundary point \hat{f}^{*+} (cf. Section 4.2). As shown in Eq. (17), this boundary point is maximal at $\hat{f}_\pi^{(0)}(t_{\pi 0} + P)$, where $\hat{f}_\pi^{(0)}$ is the rank-0 trajectory function for an arbitrary path π and is anchored at $\hat{f}_\pi^{(0)}(t_{\pi 0}) = C/P$. For $r \neq 1$ and $r = 1$, this trajectory function

⁵ In terms of latency, $(C - s)/C$, i.e., empty buffers, would even be preferable to higher values of ϵ . This latency effect could be captured by an additional axiom, which we do not introduce in this work.

⁶ Zarchy et al. [21] formalize fairness with the ratio of the smallest to the largest congestion-window size in the steady state. Given path selection, this ratio is always potentially 0, e.g., if an agent migrates in every time step.

⁷ To be precise, the asymptotic convergence to $\hat{f}^{(P-1)}$ permits that $\hat{f}_\pi(t)$ for $\text{rank}(\pi, t) = P - 1$ is consistently below $\hat{f}^{(P-1)}$. However, since this shortfall is infinitesimal and flow volumes converge exponentially to their equilibrium value, we treat the equilibrium as completely reached instead of only asymptotically approached.




Fig. 6. Computation of variance in congestion-window size according to the lossless Markov process in Eq. (27).

is given by Eqs. (14) and (15), respectively. In summary, the maximum loss rate is

$$\lambda(\text{MPCC}(\alpha, \beta, m, r)) \leq \begin{cases} 0 & \text{if } \hat{f}^{(0)} \leq C/P \\ q(m, r, P) \cdot (1 - m)^{P-1} - 1 + & \text{if } \hat{f}^{(0)} > C/P \\ (q(m, r, P) \cdot \sum_{p=0}^{P-2} \hat{\alpha}^{(p)} + \hat{\alpha}^{(P-1)}) \cdot \hat{\alpha}^{(P-1)} / (C/P) & \wedge r \neq 1 \\ ((1 - m)^{1-P} \cdot \sum_{p=0}^{P-2} \hat{\alpha}^{(p)} + \hat{\alpha}^{(P-1)}) \cdot \hat{\alpha}^{(P-1)} / (C/P) & \text{otherwise,} \end{cases} \quad (24)$$

where we use the abbreviation $q(m, r, P) := (1 + m \cdot r \cdot z(m, P))$.

Convergence (Axiom 3). If the network is in a lossless equilibrium ($\hat{f}^{(0)} \leq C/P$), the convergence behavior of the flow dynamics can be derived from the boundaries $\hat{f}^{(0)}$ and $\hat{f}^{(P-1)}$ of the hypothetical flow equilibrium. Given a lossy equilibrium, we can build on the range between the upper boundary point \hat{f}^{*+} and the lower boundary point \hat{f}^{*-} , for which we have derived an upper and a lower bound, respectively. From these ranges, the derivation of the convergence indicator γ and the maximum level \hat{f}_π^+ is straightforward:

$$\gamma(\text{MPCC}(\alpha, \beta, m, r)) \geq \begin{cases} \hat{f}^{(P-1)} / \hat{f}^{(0)} & \text{if } \hat{f}^{(0)} \leq C/P \\ \beta \cdot (1 - m)^{P-1} / (\lambda(\text{MPCC}(\alpha, \beta, m, r)) + 1) & \text{otherwise} \end{cases} \quad (25)$$

Fairness (Axiom 4). We consider the variance of congestion-window sizes in the equilibrium as a metric for the fairness of an MPCC algorithm:

$$\text{Var}_{i \in A}[cwnd_i(t)] = \mathbb{E}_{i \in A}[cwnd_i(t)^2] - \mathbb{E}_{i \in A}[cwnd_i(t)]^2 \quad (26)$$

As the congestion-window evolution of a single agent is a probabilistic process where any state transition only depends on the current state, we approximate $cwnd_i(t)$ for the case of lossless equilibria by means of the following Markov process with two state variables:

$$\begin{aligned} \text{if } \tau_i(t) \neq P - 1: \quad & \tau_i(t+1), cwnd_i(t+1) = \begin{cases} 0, r \cdot cwnd_i(t) & \text{prob. } m \\ \tau_i(t) + 1, cwnd_i(t) + \alpha(\tau_i(t)) & \text{prob. } 1 - m \end{cases} \\ \text{else:} \quad & \tau_i(t+1), cwnd_i(t+1) = \tau_i(t) + 1, cwnd_i(t) + \alpha(\tau_i(t)) \end{aligned} \quad (27)$$

where the initial state is given by $\tau_i(0) = cwnd_i(0) = 0$.

Computationally tractable computation of the congestion-window size variance can be done by averaging many simulation samples of the Markov process from Eq. (27), which has only linear complexity in t and yields the expectation of the congestion-window size by the central limit theorem. Fig. 6 illustrates that the variance of $cwnd_i(t)$ has a limit for $t \rightarrow \infty$.

Regarding lossy equilibria, the Markov process from Eq. (27) must be adapted as shown in Fig. 7. In particular, we assume that every path encounters loss with probability p_ℓ in any time step, except if the path has experienced loss in the previous time step (as there are no consecutive loss events on the same path in the lossy equilibria in Section 4). If the agent is using a lossy path, but does not leave the path, the congestion-window size is multiplicatively decreased as shown in transition *Decrease* in Fig. 7. Like for lossless equilibria, a simulation-based approach enables to efficiently compute the variance in congestion-window size (cf. Fig. D.15 in the appendix). This figure suggests that the variance limit




Fig. 7. Markov process for congestion-window size given lossy equilibria (Notation: $\bar{p} = 1 - p$). Moreover, $p_{I \rightarrow I} = 1$ if $\tau_i(t) \bmod P \neq (P - 1)$ then $\bar{p}_\ell \cdot \bar{m}$ else \bar{p}_ℓ , $p_{I \rightarrow M} = 1$ if $\tau_i(t) \bmod P \neq (P - 1)$ then m else 0 and $p_{I \rightarrow D} = 1$ if $\tau_i(t) \bmod P \neq (P - 1)$ then $p_\ell \cdot \bar{m}$ else p_ℓ .

for lossy equilibria is decreasing in loss probability p_ℓ . Moreover, since the lossy Markov process in Fig. 7 is equivalent to the lossless Markov process in Eq. (27) for $p_\ell = 0$, the variance of the lossless Markov process represents an upper bound on the variance of the lossy Markov process. Therefore, we henceforth exclusively rely on the lossless Markov process.

6. Axiom-based insights

In this section, we derive fundamental insights into the nature of end-host path selection on the basis of the axioms presented in the previous section. First, we investigate in Section 6.1 how the performance characteristics of a network change if end-host path selection is introduced. Second, we show in Section 6.2 that there are fundamental trade-offs when applying end-host path selection.

6.1. Performance effects of introducing end-host path selection

6.1.1. Evaluation method

In order to analyze how end-host path selection affects the performance characteristics of a network, we use a comparative approach: First, we characterize the performance of a network without end-host path selection based on the axioms from Section 5.1 (henceforth: Scenario (I)). Afterwards, we compare the axiomatic ratings of the network without path selection to the axiomatic characterization of the MPCC equilibria (cf. Section 5.2) that arise in the same network given end-host path selection (henceforth: Scenario (II)).

We base the comparison on a network with N agents and a total bottleneck capacity C distributed over P paths with equal bottleneck capacity C/P . All agents adopt the same CC protocol $CC_i(\alpha, \beta)$ (cf. Eq. (3)) in Scenario (I), whereas they employ a multi-path version $MPCC_i(\alpha, \beta, m, r)$ of this CC protocol in Scenario (II). Moreover, while the agent distribution on paths is dynamically determined in Scenario (II), the agent distribution in Scenario (I) is static: To identify the worst-case effects of end-host path selection, let this static agent distribution be optimal, i.e., $a_\pi = N/P$.

In the following, we rate the CC protocol $CC_i(\alpha, \beta)$ for Scenario (I) with respect to the axioms and perform a comparison with Scenario (II). Moreover, we both quantify and interpret the changes in the axiom metrics that are due to the introduction of end-host path selection. These changes are also visualized in Figs. 8 and 9: For any m and every equilibrium class (lossless or lossy), the possible range of the metric change is shown for two different additive-increase functions and an example network. We distinguish a constant additive-increase function $\alpha_1(\tau) = 1$ and an additive-increase function α_S in the style of TCP Slow Start: $\alpha_S(\tau) = 2^\tau$ if $\tau < 5$ and $\alpha_S(\tau) = 1$ otherwise. The range associated with each value of m is $[\min_{r \in R(m)} \Delta(m, r), \max_{r \in R(m)} \Delta(m, r)]$, where Δ is the difference metric as a function of m and r and $R(m)$ contains all values of r that produce a valid equilibrium of the given class (lossless or lossy) in the example network given m .

6.1.2. Efficiency (Axiom 1).

Given that the employed protocol $CC_i(\alpha, \beta)$ eventually exhausts the capacity of any path, the efficiency level is given by the lowest possible flow volume that results from loss. This lower bound is determined by the multiplicative decrease β applied to a flow volume that is infinitesimally above the capacity limit:

$$\forall \pi \in \Pi. \epsilon(CC_i(\alpha, \beta)) = \frac{\beta \cdot C/P}{C/P} = \beta \quad (28)$$

We now compare this efficiency level to the MPCC efficiency levels from Eq. (23) and analyze the efficiency change $\Delta \epsilon = \epsilon(MPCC(\alpha, \beta, m, r)) - \epsilon(CC(\alpha, \beta))$ that is due to the introduction of end-host path selection. For a visualization of this efficiency change, consider Fig. 8(a).




Fig. 8. Effects of end-host path selection for an example network with $P = 3$, $\beta = 0.7$, $N = 1000$, and $C/P = 12000$.

If the efficiency level of the MPCC dynamics is determined by a lossless equilibrium, then $\Delta\epsilon$ is given by $P \cdot \hat{f}^{(P-1)} / C - \beta$. As $\hat{f}^{(P-1)}$ is a decreasing function of the migration rate m and an increasing function of the reset softness r , end-host path selection is more likely to negatively affect ϵ for high migration rates and hard resets on path switch:

Insight 4 (Efficiency Effects of Path Migration and Resets in Lossless Equilibria). *The more likely agents are to migrate away from a path at any single point in time, the further the bottleneck-link utilization can drop, and if agents start out with a small congestion window every time they switch to a new path, utilization (and therefore efficiency) are relatively low.*

Nonetheless, it is possible that the introduction of end-host path selection leads to a higher level of efficiency. The computations for the example network, visualized in Fig. 8(a), show that for low values of m and high values of r , introducing end-host path selection can increase efficiency.

In contrast, if the MPCC efficiency level is determined by a lossy equilibrium, then $\Delta\epsilon$ is given by $\beta \cdot (1-m)^{P-1} - \beta$, which is bound to be negative. This fact allows the interpretation that end-host path selection strictly lowers the efficiency in case of loss, as emigration from a path reinforces the utilization plunge created by the CC loss reaction, i.e., the multiplicative decrease β . As Fig. 8(a) shows, such less efficient lossy equilibria are bound to exist for low values of m , for which there is no value of r such that a lossless equilibrium can arise. This insight points to a subtle relationship between migration rates and efficiency:

Insight 5 (Inefficient Equilibria due to Low Migration). *While lowering the migration rate can increase the efficiency of end-host path selection, very low migration rates necessarily lead to inefficient (lossy) equilibria, which make end-host path selection detrimental to efficiency compared to a scenario without path selection.*

6.1.3. Loss avoidance (Axiom 2).

In Scenario I, the worst-case loss rate occurs if flow f_π on path π is exactly at the capacity limit C_π , and there is an additional increase by all agents on the path:

$$\forall \pi \in \Pi. \lambda(\text{CC}_i(\alpha, \beta)) = \frac{\alpha^{\max} \cdot a_\pi}{C_\pi} = \frac{\alpha^{\max} \cdot N}{C}, \quad (29)$$

where $\alpha^{\max} = \max_{\tau \in \mathbb{N}_{>0}} \alpha(\tau)$ to represent the maximum possible loss.

In case of lossless equilibria of the MPCC dynamics, it is clear that $\Delta\lambda$ (defined analogously to $\Delta\epsilon$) is negative, i.e., the loss rate can be reduced (to 0). This improvement in $\Delta\lambda$ is shown in Fig. 8(b) for all values of m for which there is a value of r such that a lossless equilibrium arises.

If a lossy equilibrium is present, the effects of end-host path selection are more ambivalent. In that case, the maximum loss rate in the path-aware network is proportional to $\hat{f}^{(0)}$: the larger the hypothetical limit value $\hat{f}^{(0)}$ of the trajectory function, the stronger the increase of the trajectory function at level C_π and thus the higher the loss rate. As $\hat{f}^{(0)}$ is proportional to r and effectively infinite for $r = 1$, the highest loss rate for every value of m is achieved for $r = 1$, which yields the following intuitive insight:

Insight 6 (Loss Effects of Soft Resets). *If agents only perform soft resets of the congestion-window size when switching paths, this can result in high loss on the newly selected path.*

In contrast, if m and r are such that the equilibrium is only marginally lossy, i.e., $\hat{f}^{(0)}$ is only infinitesimally larger than C_π , then the maximum loss rate in a lossy equilibrium is arbitrarily close to 0 (similar to a lossless equilibrium). However, a value of r that achieves $\hat{f}^{(0)} \approx C_\pi$ may not exist given a (low) value of m , in which case the reduction of the loss rate to 0 is not possible. Therefore, we arrive at a counter-intuitive insight that mirrors Insight 5:




Fig. 9. Illustration of effects of end-host path selection on the basis of the same example network as in Fig. 8.

Insight 7 (Loss Effects of Low Migration). Loss is not minimized by minimizing the migration rate m , as low migration rates may prohibit the emergence of completely lossless equilibria.

Furthermore, Fig. 8(b) allows another non-obvious insight:

Insight 8 (Loss Effects of Path Selection with Variable Additive-Increase Functions). The benefits of end-host path selection in terms of loss are particularly large if additive-increase functions with high inherent variability (such as α_S in Fig. 8) are used by the agents.

In that case, end-host path selection may reduce loss because it leads to de-synchronization of the continuity time τ between agents: If all agents tend to have the same continuity time τ , there is a chance that many agents have continuity time τ_{\max} with $\alpha(\tau_{\max}) = \alpha^{\max}$ at the same time, resulting in high loss. In contrast, agent migration due to path selection causes more heterogeneity in τ and therefore leads to an averaging of $\alpha(\tau)$, which reduces the aggregate additive increase and therefore the maximum possible loss. While this observation may first seem like an unfair comparison of a maximum to an average, the averaging of additive increases is exactly the fundamental feature of path selection that reduces the possible maximum of aggregate additive increase compared to a scenario without path selection.

6.1.4. Convergence (Axiom 3).

The convergence level γ is determined by the minimum and the maximum possible flow volume, as derived above:

$$\gamma(\text{CC}_i(\alpha, \beta)) = \frac{\beta \cdot C}{C + \alpha^{\max} \cdot N} \quad (30)$$

In the case of lossless equilibria, end-host path selection can increase stability if

$$\frac{\hat{f}^{(P-1)}}{\hat{f}^{(0)}} = \frac{\sum_{p=0}^{P-2} \hat{\alpha}^{(p)} + \hat{\alpha}^{(P-1)} \cdot (1-m)^{P-1}}{(1+m \cdot r \cdot z(m, P)) \cdot \sum_{p=0}^{P-2} \hat{\alpha}^{(p)} + \hat{\alpha}^{(P-1)}} > \gamma(\text{CC}_i(\alpha, \beta)), \quad (31)$$

which is unsurprisingly true for a low migration rate m and hard resets $r \approx 0$. However, analogously to efficiency and loss, convergence surprisingly suffers from very low migration rates m , as this causes lossy equilibria, which are inferior to lossless equilibria in terms of convergence (cf. Fig. 9(a)).

The convergence in these lossy equilibria benefits from low migration rates and hard resets, without the exception for very low migration rates that exists for lossless equilibria. While such lossy equilibria might be inferior to lossless equilibria in terms of convergence, lossy equilibria of end-host path selection might still be preferable to no end-host path selection at all, as Fig. 9(a) suggests for the lossy equilibria for α_S . Similar to Insight 8, the reason for this improvement is the de-synchronization of the continuity time brought about by agent migration, which reduces the variance of the aggregate additive increase and thus the flow-volume fluctuations. Contrary to the widespread belief that end-host path selection necessarily hurts stability (in the sense of the convergence axiom), our analysis thus shows that network stability can in fact benefit from end-host path selection.

6.1.5. Fairness (Axiom 4).

Given simultaneous sending start and no path selection, perfect synchronization implies that all agents always have exactly the same congestion-window size, i.e., $\eta = 0$. Moreover, Zarchy et al. show that even if some agents start sending

after others, CC protocols generally tend to come close to perfect fairness [21]. To find the worst-case effects of end-host path selection, we thus assume perfect fairness in the scenario without path selection:

$$\eta(\text{CC}_i(\alpha, \beta)) = 0 \quad (32)$$

Hence, the fairness change $\Delta\eta$ due to end-host path selection is equal to the fairness level η of the MPCC dynamics, which has been computed as a function of the migration rate m in Fig. 9(b) for two different additive-increase functions. In Fig. 9(b), the lowest values for η , i.e., the highest fairness, is achieved for very high migration rates $m \approx 1$, which leads to the following insight:

Insight 9 (Fairness Effects of Path Migration). *In a system with end-host path-selection, a very high migration rate m leads to optimal fairness.*

This phenomenon can be intuitively explained as follows: If the migration rate is high, any agent is likely to reset its congestion-window size in any time step, which results in a compact distribution of the congestion-window size. Under a low migration rate, some agents may reach a high congestion-window size due to uninterrupted growth, while a few agents per time step perform a reset, which leads to a high variance of the congestion-window size distribution.

The effects of the reset softness r on η are more nuanced. As Fig. 9(b) shows, the fairness metric η is generally higher for lossy equilibria, which appear for high reset softness, than for lossless equilibria, i.e., lossless equilibria are fairer. However, as mentioned in Section 5.2, the fairness metric for lossy equilibria is computed for infinitesimal loss probability p_ℓ ; for any higher p_ℓ , η is lower, which complicates a comparison to lossless equilibria. Also, for a low migration rate m , lossy equilibria with a high reset softness are associated with lower η than lossless equilibria. The reason behind this phenomenon is that soft resets reduce the difference in congestion-window size between the agents that have not migrated in a long time (and therefore have a large congestion-window) and the agents that have just recently migrated and reset their congestion-window size.

Finally, while end-host path selection seems to reduce fairness as captured by η , we note that η only represents the fairness at any single point in time. However, under low migration rates, there may still be very high *inter-temporal* fairness. If the migration probability is low, any agent has a high probability to uninterruptedly grow its congestion window for a long time. If the congestion-window sizes of any agent were averaged over a certain time span, the distribution of such average congestion-window sizes would have low variance. We leave this more complex fairness analysis as an interesting task for future work.

6.2. Fundamental trade-offs

In Sections 5.2 and 6.1, the dependency of the MPCC dynamics on the migration rate (or responsiveness) m and the reset softness r has been qualified and quantified. These characterizations allow to observe the following trade-off in the design of systems with end-host path selection:

Insight 10 (Trade-Off Regarding Migration Rate). *Efficiency ϵ and convergence γ are more favorable under **low** migration rates, whereas fairness η and responsiveness m are more favorable under **high** migration rates, implying a fundamental trade-off between these axioms.*

However, we note that this trade-off is only valid *within* equilibrium classes, e.g., for comparing lossless equilibria among each other, but not *across* equilibrium classes: Lowering the migration rate below a certain (low) level restricts the set of possible equilibria to lossy equilibria, which are worse in terms of efficiency and convergence than lossless equilibria (cf. Insight 5).

Regarding loss avoidance, the effect of migration rates depends on the remaining network parameters. If resets are hard ($r \approx 0$), higher migration rates are associated with lower loss rate (as higher migration rates make lossless equilibria more likely, which are optimal in terms of loss). In contrast, if resets are soft ($r \approx 1$), lossless equilibria are impossible and the effects of the migration rate on the loss rate are unclear in general, because the migration rate non-monotonically affects the aggregate additive increase (cf. the curve for maximal $\Delta\lambda$ given α_5 and lossy equilibria in Fig. 8(b)). However, this unpredictable effect vanishes for constant additive-increase functions (such as α_1 from Section 6.1). For constant-increase functions, a higher migration rate leads to a *higher* loss rate given soft resets. This finding underlines the relevance of congestion-window adaptation on path switch: Depending on the reset softness, higher migration rates may increase or reduce loss.

Despite this subtle relationship of migration rates and the axiomatic metrics, we can identify parameters m and r that are optimal with respect to all the metrics efficiency, loss, and convergence simultaneously. These parameters are given by the lowest m such that a lossless equilibrium is still possible given a complete reset $r = 0$. These parameters yield a lossless equilibrium with high efficiency and convergence (cf. Insight 10).

Insight 11 (No Trade-Off between Efficiency, Loss Avoidance and Convergence). *Since there exist protocol parameters that are optimal with respect to efficiency, loss avoidance, and convergence simultaneously, there exists no fundamental trade-off between these metrics.*

Unfortunately, determining these optimal parameters requires knowledge about specific and variable properties of the target network, i.e., the number of agents $|A|$ and the path-bottleneck capacities C_π in the network, making it unattainable in most practical settings.

Finally, when determining the reset softness r , a further trade-off arises:

Insight 12 (*Trade-Off Regarding Reset Softness*). *There is a fundamental trade-off between convergence γ and loss λ , both of which are more favorable under **low** reset softness (hard resets), and efficiency ϵ , which is more favorable under **high** reset softness (soft resets).*

7. Related work

Traditionally, the effects of end-host path selection have been theoretically studied in the literature on *selfish routing*. In this line of research, the classic Wardrop model [14] is used to characterize stable traffic distributions (equilibria) that result from uncoordinated path selection by self-interested agents. These equilibria have been thoroughly investigated with respect to their existence [31,32], their efficiency (typically termed *Price of Anarchy* [33–36]), and their convergence properties [7,37,38]. However, the Wardrop model cannot represent congestion-control dynamics appropriately, which we consider important for characterizing the impact of end-host path selection on network performance.

In research about multi-path congestion control, there has been widespread use of fluid models which can better represent congestion-control dynamics [12,15–17]. However, also these models focus on representing equilibria in terms of approximate traffic distributions on networks and do not capture stability-relevant small-scale dynamics such as congestion-window fluctuations. More applied approaches rely on reasoning from network examples and experimental validation and have been used in the design of MPTCP algorithms such as LIA [18,19] and OLIA [20]. These approaches are rather suited for the design of concrete protocols than for the elicitation of fundamental properties of end-host path selection. Moreover, MPTCP research typically only investigates the effects of path selection by scrutinizing friendliness concerns between single-path and multi-path TCP users in the same network, not by looking at the impact that the introduction of end-host path selection has on aggregate performance based on various metrics.

In contrast, the axiomatic approach used in this paper allows to qualify and quantify the performance impact of path selection on a fundamental level while taking congestion-control dynamics into account. Thanks to this power, the axiomatic perspective has been applied to various topics beyond game theory: In computer science, for example, research on congestion control [21], routing protocols [39], and recommendation systems [40] have benefited from axiom-based approaches.

The effects of end-host path selection have also been characterized by Wang et al. [24], whose ‘cost of not splitting in routing’ captures the difference in network utility between a scenario where end-hosts select a single path and a scenario where multiple paths can be selected. However, this work differs from ours in investigating static rate allocations instead of dynamic rate evolution, in evaluating a single metric (utility) instead of multiple axioms, and in contrasting different modes of end-host path selection instead of contrasting path selection with path pre-determination.

8. Conclusion

Motivated by a stability concern about end-host path selection, we qualify and quantify the performance impact of such path-selection-induced instability in this work. More precisely, we analyze a general network in which end-hosts employ greedy load-adaptive path selection and characterize the resulting traffic pattern with respect to five metrics of interest (“axioms”): efficiency, loss avoidance, convergence, fairness and responsiveness. Through this analysis, we show how the performance impact of end-host path selection depends on the path-migration behavior, the underlying congestion-control protocol, and the structure of the network. Among the dependencies that we present and explain, there are both intuitive, well-known dependencies (e.g., high migration rates decrease efficiency) and non-intuitive, more complex dependencies (e.g., very low migration rates increase loss). Moreover, we show that there are fundamental limitations such that no multi-path congestion-control protocol can optimize all metrics simultaneously.

We understand our work as a first step, which allows many avenues for follow-up research. For example, it would be interesting to extend the model for additional congestion-control behaviors (e.g., latency-based protocols or model-based protocols such as BBR [41]), additional path-switching behaviors (e.g., based on a path-switching probability proportional to the load difference between paths) and more general networks. However, while our insights admittedly stem from a simplified model, we believe that the illustrated dependencies and the axiomatic reasoning in general can inform the discussion about the merits and perils of end-host path selection.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We gratefully acknowledge support from ETH Zurich, from SNSF for project ESCALATE (200021L_182005), and from the Austrian Science Fund (FWF) for project I 4800-N (ADVISE), 2020–2023. Moreover, we thank Joel Wanner for his helpful feedback, the anonymous reviewers for their careful reviews, and Edmundo de Souza e Silva for his shepherding.

Appendix A. Analysis of the continuity-time distribution

The agent dynamics involved in P -step oscillation (Definition 1) allow to estimate how long the agents on a path have already been using that path without a packet loss, i.e., allow to characterize the distribution of the *continuity time* introduced above. For the following analysis, we introduce the notation $\theta(\pi, t)$, which shall denote the time since the most recent loss event on path π at time t .

We now derive a probability distribution $\mathbb{P}[\tau_i(t) = \tau]$, giving the probability that agent i on path $\pi_i(t)$ has continuity time $\tau \in \mathbb{N}$ at time t . This distribution will later be used to determine the expected congestion-window increase $\hat{\alpha}_\pi(t)$ in Eq. (4b). We consider an arbitrary agent $i \in A$ at an arbitrary time t , residing on path $\pi_i(t)$. Clearly, agent i must have continuity time $\tau_i(t) = 0$ right after a loss event, i.e., whenever $\theta(\pi_i(t), t) = 0$, irrespective of the rank of $\pi_i(t)$:

$$\forall t \text{ s.t. } \theta(\pi_i(t), t) = 0, p \in [P]. \quad \mathbb{P}[\tau_i(t) = 0 \mid \text{rank}(\pi_i(t), t) = p] = 1 \quad (\text{A.1})$$

However, in the subsequent time steps, where $\theta(\pi_i(t), t) > 0$, the continuity-time distribution of agent i on path $\pi_i(t)$ depends on the rank of that path. If $\text{rank}(\pi_i(t), t) = 0$ or, equivalently, $\text{rank}(\pi_i(t-1), t-1) = P-1$, all the $\hat{a}^{(P-1)}$ agents that were on path π in the last time step $t-1$ have remained on the path and have increased their continuity time by 1, but their *relative share* is reduced by on-migration from other paths:

$$\begin{aligned} \forall t \text{ s.t. } \theta(\pi_i(t), t) > 0, \tau > 0. \\ \mathbb{P}[\tau_i(t) = \tau \mid \text{rank}(\pi_i(t), t) = 0] \\ = \mathbb{P}[\tau_i(t-1) = \tau - 1 \mid \text{rank}(\pi_i(t-1), t-1) = P-1] \cdot \hat{a}^{(P-1)} / \hat{a}^{(0)} \\ = \mathbb{P}[\tau_i(t-1) = \tau - 1 \mid \text{rank}(\pi_i(t-1), t-1) = P-1] \cdot (1-m)^{P-1} \end{aligned} \quad (\text{A.2})$$

All the $m \cdot (N - \hat{a}^{(P-1)})$ agents that migrated from the other paths have continuity time 0:

$$\forall t \text{ s.t. } \theta(\pi_i(t), t) > 0. \quad \mathbb{P}[\tau_i(t) = 0 \mid \text{rank}(\pi_i(t), t) = 0] = \frac{m \cdot (N - \hat{a}^{(P-1)})}{\hat{a}^{(0)}} = 1 - (1-m)^{P-1}. \quad (\text{A.3})$$

If path $\pi_i(t)$ has $\text{rank}(\pi_i(t), t) \neq 0$, the continuity-time distribution has been shifted up by 1 in the last time step, but is otherwise unaffected:

$$\begin{aligned} \forall t \text{ s.t. } \theta(\pi_i(t), t) > 0, p \in [P] \setminus \{0\}, \tau > 0. \\ \mathbb{P}[\tau_i(t) = \tau \mid \text{rank}(\pi_i(t), t) = p] = \mathbb{P}[\tau_i(t-1) = \tau - 1 \mid \text{rank}(\pi_i(t-1), t-1) = p-1] \end{aligned} \quad (\text{A.4})$$

These recursive characterizations of the probability distribution are equivalent to the following explicit definition of the continuity-time distribution, which is visualized in Fig. B.12:

Insight 13. At time t , the probability that an agent i on a path $\pi_i(t)$ with rank p and time since last loss $\theta(\pi_i(t), t)$ has continuity time τ is

$$\begin{aligned} \mathbb{P}(\tau; t, p) &:= \mathbb{P}[\tau_i(t) = \tau \mid \text{rank}(\pi_i(t), t) = p] \\ &= \begin{cases} (1-m)^{\lceil \frac{\tau-p}{P} \rceil (P-1)} & \text{if } \tau = \theta(\pi_i(t), t) \\ (1 - (1-m)^{P-1}) \cdot (1-m)^{\lfloor \frac{\tau}{P} \rfloor (P-1)} & \text{if } \tau < \theta(\pi_i(t), t) \wedge \tau \bmod P = p \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (\text{A.5})$$

On a path π with rank p at time t , the expected additive increase $\hat{\alpha}_\pi(t)$ at time t is therefore:

$$\hat{\alpha}_\pi(t) = (1-m)^{\lceil \frac{\theta-p}{P} \rceil (P-1)} \cdot \alpha(\theta) + \sum_{k=0}^{\lceil (\theta-p)/P \rceil - 1} (1 - (1-m)^{P-1}) \cdot (1-m)^{k(P-1)} \cdot \alpha(Pk + p) \quad (\text{A.6})$$

where $\theta = \theta(\pi, t)$.

For increasing time since the last loss ($\theta \rightarrow \infty$), the expected average additive increase on a path with rank p converges to the following quantity, which can be easily computed for any additive-increase function α :

$$\hat{\alpha}^{(p)} = \sum_{k=0}^{\infty} (1 - (1-m)^{P-1}) \cdot (1-m)^{k(P-1)} \cdot \alpha(Pk + p) \quad (\text{A.7})$$




Fig. B.10. Simulated MPCC dynamics $\{(a_\pi(t), f_\pi(t))\}_{\pi \in \Pi}$ (gray dashed) and expected MPCC dynamics $\{(\hat{a}_\pi(t), \hat{f}_\pi(t))\}_{\pi \in \Pi}$ (in color) for $P = 3$, $N = 1000$, $\alpha(\tau) = 1$, $\beta = 0.7$, and $C_\pi = 12000$ for every $\pi \in \Pi$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Appendix B. Approximation accuracy

The expected MPCC dynamics in Eq. (4) are an approximation of the actual probabilistic MPCC dynamics in Eq. (2), which are unsuitable for analytic investigation. In order to demonstrate the accuracy of this approximation, we present a comparison between the actual dynamics and the expected dynamics for a selection of parameters in Fig. B.10. In each sub-figure, the actual MPCC dynamics from Eqs. (2a) and (2b) are simulated and shown with light gray lines, and the expected dynamics are computed and drawn with colored lines (agent dynamics in dotted lines, flow dynamics in dashed lines).

In Fig. B.10(a) and Fig. B.10(b) (i.e., $r \neq 1$), the expected dynamics are compared with results from 5 simulation runs of the actual dynamics. The expected dynamics appropriately capture the structure of both the agent dynamics and the flow dynamics, in particular the curvature, the convergence behavior and the reaction to loss (e.g., at $t = 35$ in Fig. B.10(b)). As the actual dynamics are realizations of a random variable, their values deviate from the expectation; however, the variance is modest.

In Fig. B.10(c) (i.e., $r = 1$), the actual flow dynamics look more different from the expected flow dynamics than for $r \neq 1$. This difference is due to loss events at different points in time, which can even result in case of low variance, but make the dynamics look quite different. However, the pattern of recurring loss is well captured by the expected dynamics. In order to make this similarity visible, only one simulation run of the actual dynamics is shown.

The analysis above is repeated for more paths and a non-constant additive-increase function in Appendix B. In particular, we repeat this analysis for constant additive increase, but with $P = 5$ (cf. Fig. B.11(a)), as well as with an additive-increase function that mimics TCP slow-start behavior ($\alpha_s(\tau) = 2^\tau$ if $\tau < 5$ else 1) for both $P = 3$ (cf. Fig. B.11(b)) and $P = 5$ (cf. Fig. B.11(c)).

Appendix C. Logical consistency of P -step oscillation

In Section 3.3, we have shown that given P -step oscillation and without capacity limits, the flow dynamics exponentially converge to a dynamic equilibrium where the rank- p path carries flow volume $\hat{f}^{(p)}$ in every time step. The general rank- p equilibrium flow volume is given by the following term:

$$\hat{f}^{(p)} = \frac{\sum_{p'=0}^{p-1} (1-m)^{p-p'} \hat{\alpha}^{(p')} \cdot \hat{a}^{(p')} + (1-m)^p \cdot \hat{\alpha}^{(p-1)} \cdot \hat{a}^{(p-1)}}{1 - (1+m \cdot r \cdot z(m, P)) \cdot (1-m)^{p-1}} + \frac{(1+m \cdot r \cdot z(m, P)) \cdot \sum_{p'=p}^{P-2} (1-m)^{p-1+p-p'} \cdot \hat{\alpha}^{(p')} \cdot \hat{a}^{(p')}}{1 - (1+m \cdot r \cdot z(m, P)) \cdot (1-m)^{p-1}} \quad (\text{C.1})$$

Interestingly, analyzing the equilibrium flow volumes $\{\hat{f}^{(p)}\}_{p \in [P]}$ allows to draw conclusions about the occurrence of P -step oscillation for a certain parameter combination, which works by logical contraposition: If P -step oscillation occurs for a certain parameter combination, then P -step oscillation produces the equilibrium flow volumes $\{\hat{f}^{(p)}\}_{p \in [P]}$. However, if the equilibrium flow volumes are themselves inconsistent with P -step oscillation, i.e., if $\hat{f}^{(p)} < \hat{f}^{(p+1)}$ for some $p \in [P-1]$, then the equilibrium cannot exist and there is a contradiction. This contradiction suggests that P -step




Fig. B.11. Comparison of model and simulations to demonstrate approximation accuracy.

oscillation is fundamentally impossible for the given parameter combination, as P -step oscillation would have produced the equilibrium flow volumes if it had occurred.⁸

Based on this reasoning, we can find a parameter sub-space for which P -step oscillation is impossible. More precisely, given any parameter combination, we can compute the equilibrium flow volumes $\{\hat{f}^{(p)}\}_{p \in [P]}$ and check if $\hat{f}^{(p)} < \hat{f}^{(p+1)}$ for

⁸ Note that the inverse is not true: The absence of a contradiction does not mean that P -step oscillation necessarily occurs for a given parameter combination.




Fig. B.12. Continuity-time distribution for $P = 3$, $m = 0.15$, and different θ . The dashed lines represent the function $(1 - (1 - m)^{P-1}) \cdot (1 - m)^{(\tau - p)/P}$.




Fig. B.13. Visualization of parameter sub-space that is inconsistent with P -step oscillation for different additive-increase functions $\alpha_1(\tau) = 1$ and $\alpha_S(\tau) = (2^\tau \text{ if } \tau < 5 \text{ else } 1)$.

any $p \in [P - 1]$. As Eq. (C.1) shows, the parameter space for the equilibrium flow volumes consists of the migration rate m , the reset softness r , the number of paths P , the additive-increase function $\alpha(\tau)$, and the number of agents N (appearing in $\hat{a}^{(p)}$). Luckily, as N is a linear coefficient of $\hat{f}^{(p)}$ and $N > 0$, N can be eliminated in the inequality $\hat{f}^{(p)} < \hat{f}^{(p+1)}$. We performed such an exploration of the parameter space with a focus on m and r , yielding the results in Fig. B.13. These results indicate that for the two analyzed additive-increase functions, P -step oscillation is never logically inconsistent for 2 paths and only rarely logically inconsistent for higher number of paths. While not a definitive proof for the prevalence of P -step oscillation, these results suggest that the notion of P -step oscillation is a sound concept for most parameter combinations.

Appendix D. Additional figures

This appendix section contains additional figures that illustrate concepts presented in the main body of the paper. Fig. D.14 presents a simulation-based validation of the lower bounds on the flow volume in lossy equilibria, derived in Section 4.2. Fig. D.15 presents the variance in congestion-window size given a lossy equilibrium, computed from simulation of the lossy Markov process in Fig. 7 in Section 5.2.




Fig. D.14. Simulation-based validation of lower bounds on efficiency ϵ for lossy equilibria as derived in Section 4.2. Simulation parameters of interest include $C_\pi = 12000$, $N = 1000$, and $\alpha(\tau) = 1$.




Fig. D.15. Simulation-based computation of variance in congestion-window size according to the lossy Markov process in Fig. 7 for different values of loss probability p_ℓ .

References

- [1] João Luís Sobrinho, Miguel Alves Ferreira, Routing on multiple optimality criteria, in: Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, 2020, pp. 211–225.
- [2] Muhammad Shahbaz, Lalith Suresh, Jennifer Rexford, Nick Feamster, Ori Rottenstreich, Mukesh Hira, Elmo: Source routed multicast for public clouds, in: Proceedings of ACM SIGCOMM, Association for Computing Machinery (ACM), 2019, pp. 458–471.
- [3] David B. Johnson, David A. Maltz, Dynamic source routing in ad hoc wireless networks, in: Mobile Computing, Springer, 1996.
- [4] Giacomo Giuliani, Tobias Klente, Markus Legner, David Basin, Adrian Perrig, Ankit Singla, Internet backbones in space, ACM SIGCOMM Comput. Commun. Rev. (2020).
- [5] Clarence Filisfilis, Nagendra Kumar Nainar, Carlos Pignataro, Juan Camilo Cardona, Pierre Francois, The segment routing architecture, in: IEEE Global Communications Conference (GLOBECOM), 2015.
- [6] David Barrera, Laurent Chuat, Adrian Perrig, Raphael M. Reischuk, Pawel Szalachowski, The SCION internet architecture, Commun. ACM (2017).
- [7] Simon Fischer, Berthold Vöcking, Adaptive routing with stale information, Theoret. Comput. Sci. (2009).
- [8] Anees Shaikh, Jennifer Rexford, Kang G. Shin, Evaluating the impact of stale link state on quality-of-service routing, IEEE/ACM Trans. Netw. (2001).
- [9] Spencer Dawkins, Path aware networking: Obstacles to deployment (a bestiary of roads not taken), Internet-Draft draft-irtf-panrg-what-not-to-do-16, Internet Engineering Task Force, 2020, (in preparation).
- [10] Anwar Elwalid, Cheng Jin, Steven Low, Indra Widjaja, Mate: Multipath adaptive traffic engineering, Comput. Netw. (2002).
- [11] Srihari Nelakuditi, Zhi-Li Zhang, Rose P. Tsang, David Hung-Chang Du, Adaptive proportional routing: a localized QoS routing approach, IEEE/ACM Trans. Netw. (2002).
- [12] Frank Kelly, Thomas Voice, Stability of end-to-end algorithms for joint routing and rate control, ACM SIGCOMM Comput. Commun. Rev. (2005).

- [13] Simon Scherrer, Markus Legner, Adrian Perrig, Stefan Schmid, Incentivizing stable path selection in future internet architectures, in: International Symposium on Computer Performance, Modeling, Measurements and Evaluation, 2020.
- [14] John Glen Wardrop, Some theoretical aspects of road traffic research, *Proc. Inst. Civ. Eng.* (1952).
- [15] Qiuyu Peng, Anwar Walid, Steven H. Low, Multipath TCP algorithms: theory and design, *ACM SIGMETRICS Perform. Eval. Rev.* (2013).
- [16] Peter Key, Laurent Massoulié, Don Towsley, Path selection and multipath congestion control, in: IEEE International Conference on Computer Communications (INFOCOM), IEEE, 2007, pp. 143–151.
- [17] Huaizhong Han, Srinivas Shakkottai, Christopher V. Hollot, Rayadurgam Srikant, Don Towsley, Multi-path tcp: a joint congestion control and routing scheme to exploit path diversity in the internet, *IEEE/ACM Trans. Netw.* (2006).
- [18] Damon Wischik, Costin Raiciu, Adam Greenhalgh, Mark Handley, Design, implementation and evaluation of congestion control for multipath TCP, in: *NSDI*, vol. 11, 2011, pp. 8.
- [19] Costin Raiciu, Mark Handley, Damon Wischik, Coupled congestion control for multipath transport protocols, RFC 6356, Internet Engineering Task Force (IETF), 2011.
- [20] Ramin Khalili, Nicolas Gast, Miroslav Popovic, Jean-Yves Le Boudec, Mptcp is not Pareto-optimal: Performance issues and a possible solution, *IEEE/ACM Trans. Netw.* (2013).
- [21] Doron Zarchy, Radhika Mittal, Michael Schapira, Scott Shenker, Axiomatizing congestion control, in: Proceedings of the ACM Conference on Measurement and Analysis of Computing Systems (SIGMETRICS), vol. 3 (2) pp. 1–33.
- [22] Douglas H. Blair, Robert A. Pollak, Acyclic collective choice rules, *Econometrica* (1982).
- [23] John F. Nash Jr., The bargaining problem, *Econometrica* (1950).
- [24] Meng Wang, Chee Wei Tan, Weiyu Xu, Ao Tang, Cost of not splitting in routing: Characterization and estimation, *IEEE/ACM Trans. Netw.* 19 (6) (2011) 1849–1859.
- [25] Jeonghoon Mo, Richard J. La, Venkat Anantharam, Jean Walrand, Analysis and comparison of TCP Reno and Vegas, in: INFOCOM, IEEE, 1999.
- [26] Sangtae Ha, Injong Rhee, Lisong Xu, CUBIC: a new TCP-friendly high-speed TCP variant, *ACM SIGOPS Rev.* (2008).
- [27] Wright Stevens, et al., Tcp slow start, congestion avoidance, fast retransmit, and fast recovery algorithms, RFC 2001, Internet Engineering Task Force (IETF), 1997.
- [28] Eitan Altman, K.E. Avrachenkov, B.J. Prabhu, Fairness in MIMD congestion control algorithms, *Telecommun. Syst.* (2005).
- [29] Oded Galor, *Discrete Dynamical Systems*, Springer Science & Business Media, 2007.
- [30] Aditya Akella, Srinivasan Seshan, Richard Karp, Scott Shenker, Christos Papadimitriou, Selfish behavior and stability of the internet: A game-theoretic analysis of TCP, *ACM SIGCOMM Comput. Commun. Rev.* 32 (4) (2002) 117–130.
- [31] Robert W. Rosenthal, A class of games possessing pure-strategy Nash equilibria, *Internat. J. Game Theory* (1973).
- [32] Tim Roughgarden, Routing games, in: *Algorithmic Game Theory*, vol. 18, Cambridge University Press, 2007, pp. 459–484.
- [33] Elias Koutsoupias, Christos Papadimitriou, Worst-case equilibria, in: *Annual Symposium on Theoretical Aspects of Computer Science*, Springer, 1999, pp. 404–413.
- [34] Tim Roughgarden, Éva Tardos, How bad is selfish routing? *J. ACM* (2002).
- [35] Tim Roughgarden, The price of anarchy is independent of the network topology, *J. Comput. System Sci.* (2003).
- [36] Lili Qiu, Yang Richard Yang, Yin Zhang, Scott Shenker, On selfish routing in Internet-like environments, in: Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, 2003.
- [37] William H. Sandholm, Potential games with continuous player sets, *J. Econom. Theory* (2001).
- [38] Simon Fischer, Berthold Vöcking, On the evolution of selfish routing, in: *European Symposium on Algorithms*, Springer, 2004, pp. 323–334.
- [39] Omer Lev, Moshe Tennenholtz, Aviv Zohar, An axiomatic approach to routing, in: INFOCOM, 2015.
- [40] Reid Andersen, Christian Borgs, Jennifer Chayes, Uriel Feige, Abraham Flaxman, Adam Kalai, Vahab Mirrokni, Moshe Tennenholtz, Trust-based recommendation systems: an axiomatic approach, in: Proceedings of the International Conference on World Wide Web, 2008, pp. 199–208.
- [41] Neal Cardwell, Yuchung Cheng, C. Stephen Gunn, Soheil Hassas Yeganeh, Van Jacobson, BBR: Congestion-based congestion control, *Queue* 14 (5) (2016) 20–53.



Simon Scherrer is a doctoral student in the Network Security Group at ETH Zurich, Switzerland, under the supervision of Professor Adrian Perrig. His research interests include game theory in networks, the investigation of network architectures with end-host path selection, and probabilistic traffic monitoring. He obtained an MSc degree from ETH Zurich in 2019.



Markus Legner is a postdoctoral researcher in the Network Security Group headed by Prof. Dr. Adrian Perrig at ETH Zurich, Switzerland. He is working on SCION, a next-generation path-aware Internet architecture that at the same time provides high security and high availability. In particular, he focuses on improving the control-plane security and investigating effects of end-host-based path selection on network properties. Before joining the Network Security Group, Markus received his Ph.D. in theoretical physics from ETH Zurich in 2016.

Adrian Perrig is a Professor at the Department of Computer Science at ETH Zürich, Switzerland, where he leads the network security group. He earned his MS and PhD degrees in Computer Science from Carnegie Mellon University, and spent three years during his PhD at the University of California at Berkeley. He received his BSc degree in Computer Engineering from EPFL.



Stefan Schmid is a Full Professor at TU Berlin, Germany, Fraunhofer SIT, Germany, and University of Vienna, Austria. He completed his MSc and PhD at ETH Zurich, and he has worked as a Postdoc at TU Munich and University of Paderborn, as a Senior Research Scientist at T-Labs in Berlin, and as an Associate Professor at Aalborg University, Denmark. Stefan Schmid received the IEEE Communications Society ITC Early Career Award 2016 and an ERC Consolidator Grant 2019.