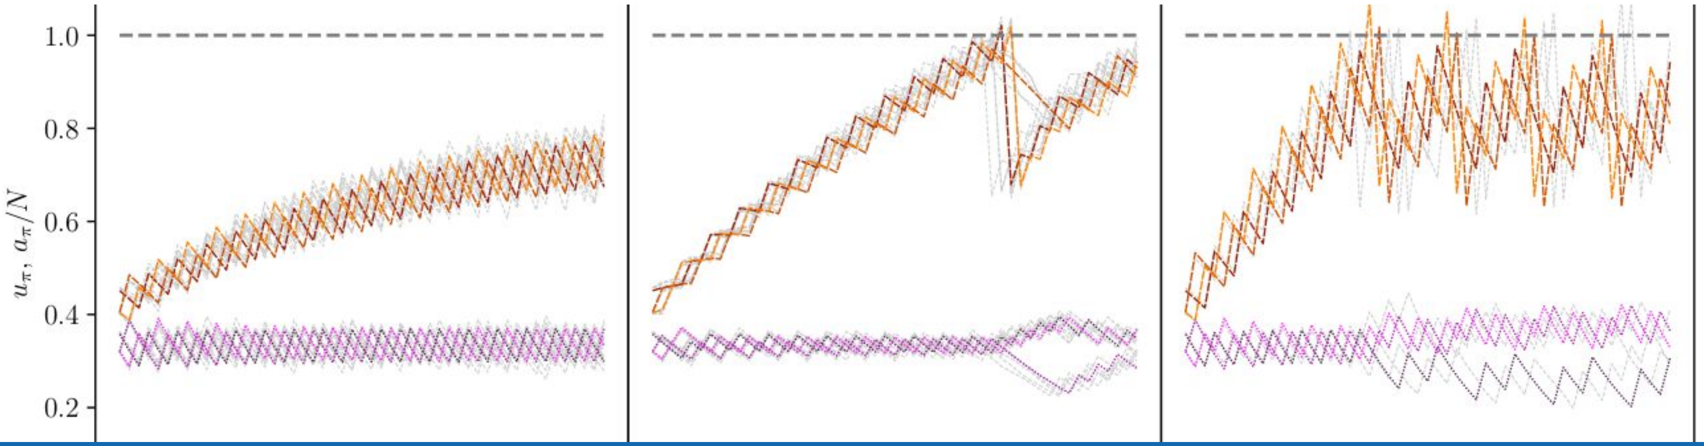

An Axiomatic Perspective on the Performance Effects of End-Host Path Selection

Simon Scherrer
ETH Zurich

Markus Legner
ETH Zurich

Adrian Perrig
ETH Zurich

Stefan Schmid
TU Berlin &
University of Vienna

PERFORMANCE 2021
Virtual Conference

Motivating example: Lane splitting

(wikimedia.org)

- Banned in many places because of intuitive safety concerns

However:

- Research fails to confirm lane splitting as hazardous
- Lane-splitting bans reduce traffic speed

→ Lane-splitting bans may be more costly than lane-splitting itself

(M. Sperley, Motorcycle Lane-Sharing: Literature Review, 2010)

General insight: If problem is never properly characterized, solution may be worse than the problem

Parallel to end-host path selection:

- Widespread intuition that unregulated end-host path selection causes damage through oscillation
- Many proposals for path-selection regulation systems which introduce overhead
- Little rigorous research on extent and effects of oscillation problem

Research question:
How exactly and *by how much* does oscillatory path selection decrease network performance?

Research approach:
Inspiration from **axiomatic approach** by Zarchy et al. *(SIGMETRICS'20)*

➔ Previous analytical approaches in congestion-control research may be suitable for **analysis or design of specific protocols**…

… but not for the discovery of **fundamental constraints of the design space** (e.g., achievable optima and trade-offs with respect to different desirable properties)

axioms

33

Axiomatizing Congestion Control

DORON ZARCHY, Hebrew University of Jerusalem
RADHIKA MITTAL, University of Illinois at Urbana-Champaign
MICHAEL SCHAPIRA, Hebrew University of Jerusalem
SCOTT SHENKER, UC Berkeley, ICSI

The overwhelmingly large design space of congestion control protocols, along with the increasingly diverse range of application environments, makes evaluating such protocols a daunting task. Simulation and experiments are very helpful in evaluating the performance of designs in *specific* contexts, but give limited insight into the more general properties of these schemes and provide no information about the *inherent* limits of congestion control designs (such as, which properties are simultaneously achievable and which are mutually exclusive). In contrast, traditional theoretical approaches are typically focused on the design of protocols that achieve to specific, predetermined objectives (e.g., network utility maximization), or the analysis of specific protocols (e.g., from control-theoretic perspectives), as opposed to the inherent tensions/derivations between desired properties.

To complement today's prevalent experimental and theoretical approaches, we put forth a novel principled framework for reasoning about congestion control protocols, which is inspired by the axiomatic approach from social choice theory and game theory. We consider several natural requirements ("axioms") from congestion control protocols – e.g., efficient resource-utilization, loss-avoidance, fairness, stability, and TCP-friendliness – and investigate which combinations of these can be achieved within a single design. Thus, our framework allows us to investigate the fundamental tradeoffs between desiderata, and to identify where existing and new congestion control architectures fit within the space of possible outcomes.

CCS Concepts: • **Networks** → Network protocols;

ACM Reference Format:
Doron Zarchy, Radhika Mittal, Michael Schapira, and Scott Shenker. 2019. Axiomatizing Congestion Control. *Proc. ACM Meas. Anal. Comput. Syst.* 3, 2, Article 33 (June 2019), 33 pages. https://doi.org/10.1145/3326148

ETH zürich

Axiomatic approach: A simple example

1) Represent protocol behavior in single-bottleneck discrete-time model
2) Assess protocol behavior w.r.t. axiom metrics
3) Identify fundamental constraints of the protocol behavior

$$TCPReno(t, cwnd_i(t)) = \begin{cases} cwnd_i(t) + 1 & \text{if } f(t) \leq C \\ 0.5 \cdot cwnd_i(t) & \text{otherwise} \end{cases}$$

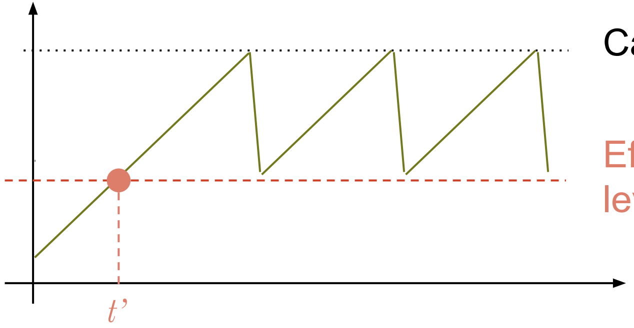

α-efficiency: $\exists\ t'.\ \ \forall t' > t.\ \ \Sigma_i\ cwnd_i(t) \geq \alpha \boldsymbol{C}$

Fundamental constraint:
TCP Reno is 0.5-efficient.

THEOREM 3. *Any loss-based protocol that is α-fast-utilizing, β-efficient, and ϵ-robust, for $\epsilon > 0$, is at most*
$\frac{3(1-\beta)}{(4(\frac{C+\tau}{1-\epsilon})-\alpha)(1+\beta)}$*-TCP friendly.*

(Zarchy et al.)

Axiomatic perspective on end-host path selection

1) We extend Zarchy et al.'s axiomatic CC approach to joint congestion control and path selection (MPCC)
   → *Multiple bottleneck links available for selection*

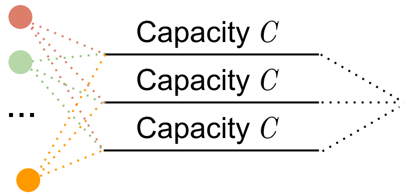

2) We are interested in the worst-case effects of oscillatory path selection
   → *Greedy path selection: Switch to least utilized path with given probability m*
   → *Sequential multi-path: Every sender uses only one path at a given moment*
   → *All paths have same capacity s.t. differences in load matter for attractiveness*
   → *Senders have same RTT, so their reactions are strongly correlated*

3) We test different policies for handling path switch in CC algorithm
   → *Reset r of congestion window upon path switch*
   → *Function $\alpha(t)$ regulating congestion-window growth after path switch*

Characterization of path selection (Part 1)

1) LLN-based approximation of stochastic process by **expected dynamics**

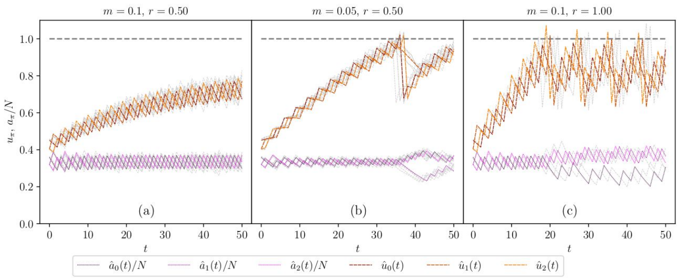

2) Identification of common **oscillation pattern** (P-step oscillation)

3) Derivation of **dynamic equilibria** produced by oscillation pattern

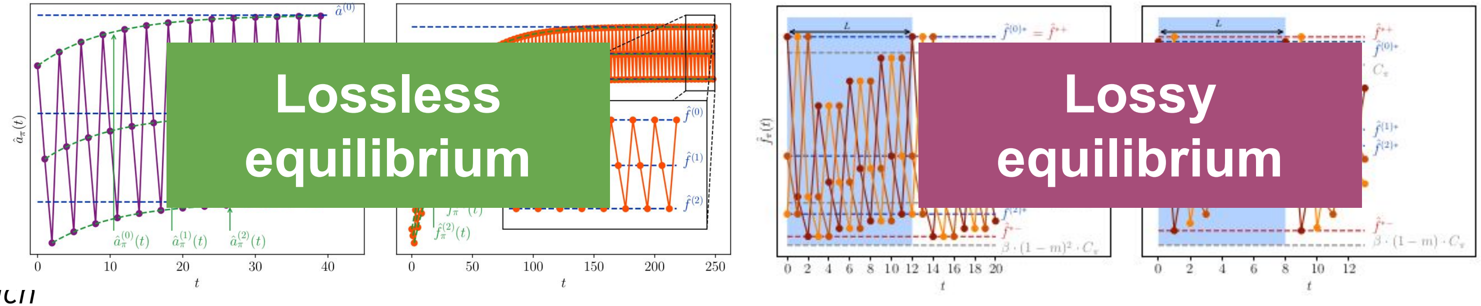

Characterization of path selection (Part 2)

4) Rating of MPCC equilibria according to **axiom metrics**

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

(Lossless equilibrium)
(Lossy equlibrium)

5) Rating of underlying CC protocol **in absence of path selection**

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

6) **Comparison** of multi-path axiom scores and single-path axiom scores

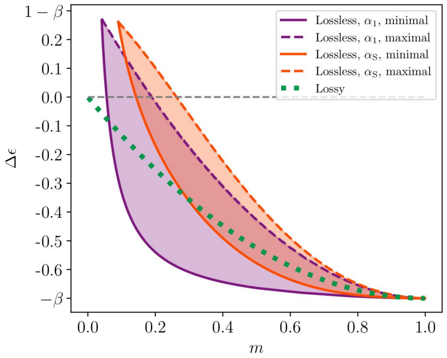

ETH zürich

Characterization of path selection (Part 3)

7) Elicitation of insights from comparison-based analysis
For example:

> *A high path-migration rate harms efficiency.*
> *However, a moderate migration rate is preferable to a very low migration rate.*

For very low *m*, only inefficient lossy equilibria are possible

Efficiency change due to path selection

Efficiency changes achievable for varying resets *r* given *m*

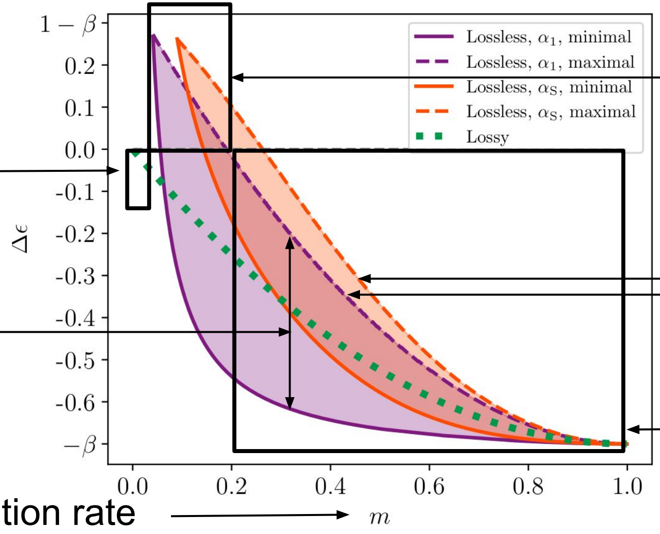

Migration rate

Modest migration rates allow to improve efficiency compared to no path selection

Different policies for cwnd growth after path switch

High migration rates make path selection *forcibly* worse than no path selection

ETH zürich

Characterization of path selection (Part 3)

7) Elicitation of insights from comparison-based analysis
For example:

> *A high path-migration rate is associated with high fairness.*

Cwnd variance as fairness metric

$$\exists t' > 0. \quad \forall t > t'. \quad \underset{i \in A}{\mathrm{Var}}\left[cwnd_i(t)\right] \leq \eta$$

Scenario without path selection:
Perfect fairness (Zarchy et al.)

$$\eta\left(CC_i(\alpha, \beta)\right) = 0$$

Scenario with path selection:
Markov process

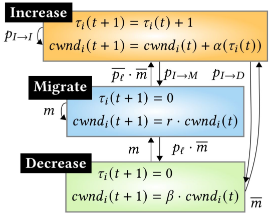

(Un)fairness change due to path selection (Low *Δη* = high fairness)

Best fairness for highest migration rate

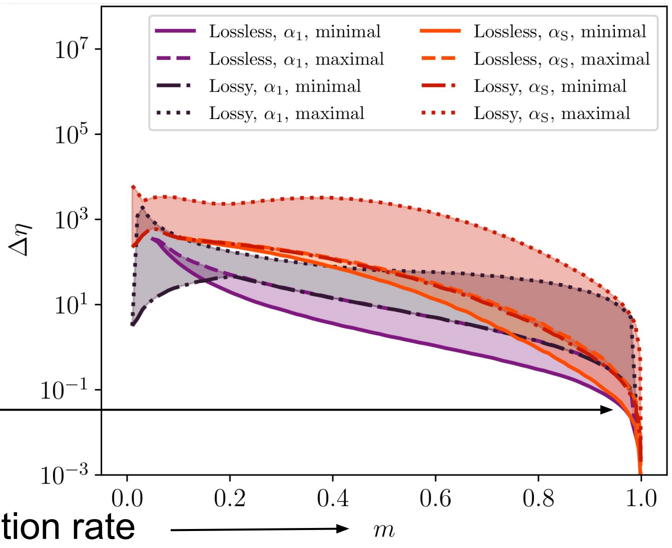

Migration rate

Characterization of path selection (Part 3)

7) Elicitation of insights from comparison-based analysis
For example

> *A high path-migration rate harms efficiency, loss avoidance and stability, but improves fairness. However, even for metrics that benefit from low migration, a moderate migration rate is preferable to a very low migration rate.*

> *Path selection may worsen efficiency because the utilization plunge created by loss is compounded with out-migration from the loss-experiencing path.*

> *A hard reset of the congestion window upon path switch improves network stability, but harms efficiency.*

etc.

Conclusion

Questions?

How exactly and *by how much* does oscillatory path selection decrease network performance?

Axiomatic approach → Multi-path extension → Equilibrium characterization → Comparative analysis

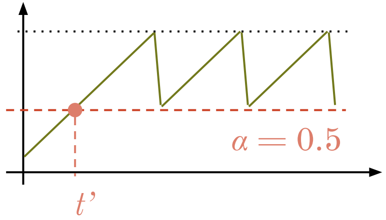

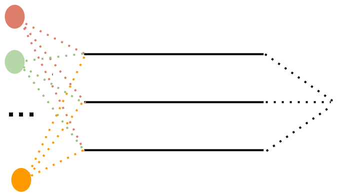

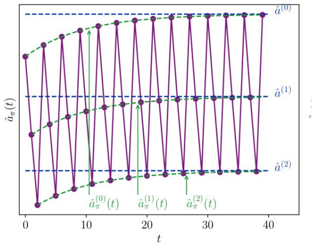

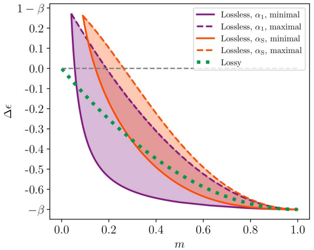

Central finding:

Oscillation from path selection is not necessarily harmful!

For moderate migration rates, it is even preferable to no path selection at all.