A Wavelet-Based Approach to Detect Shared Congestion

Min Sik Kim
The University of Texas at Austin

Coauthors: Taekhyun Kim, YongJune Shin,
Simon S. Lam, Edward J. Powers

Cooperative Congestion Control

- Better utilization of network resources
- Applications
  - Congestion Manager, path diversity
  - Improving overlay network topology
    - end system multicast, overlay routing, ...

- *Identify flows sharing a bottleneck!*

Previous Approaches to Detect Shared Congestion

- Loss-based techniques
  - Work with lossy links, drop-tail queues
  - Do *not* work with low loss rate, RED
- Delay-based techniques
  - More robust than loss-based ones
- Limitation
  - Require a *common endpoint*

Outline

- ☑ Introduction
- ☐ Basic technique
- ☐ Limitations of the basic technique
- ☐ DCW: Delay Correlation with Wavelet denoising
- ☐ Experimental results
- ☐ Summary

Model

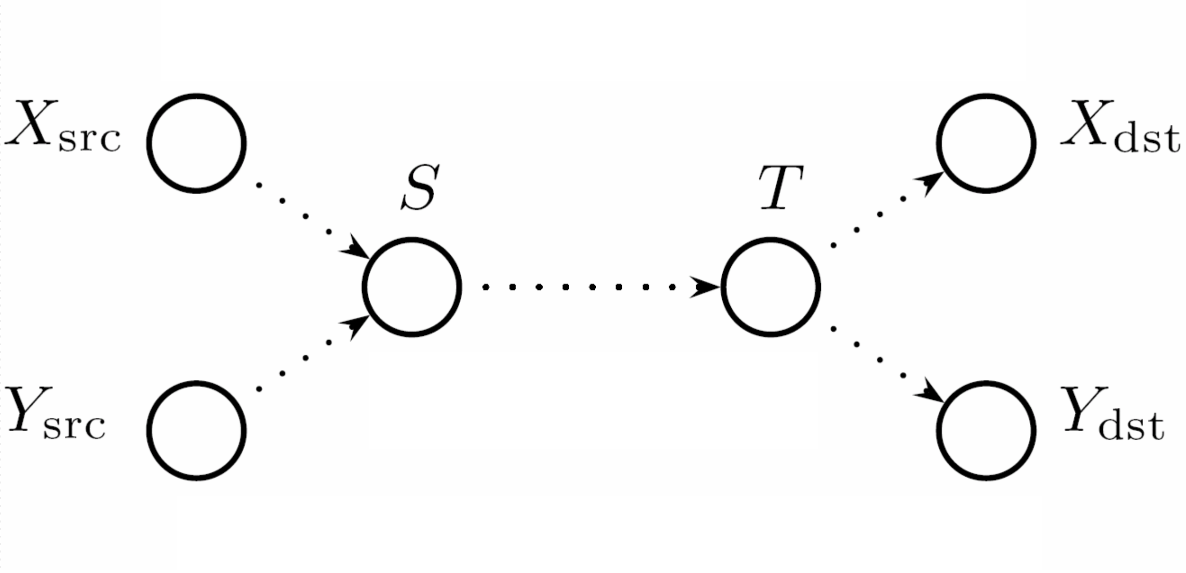

- Observations on queueing delay
  - Congested link: large fluctuations
  - Non-congested link: stable

Basic Technique

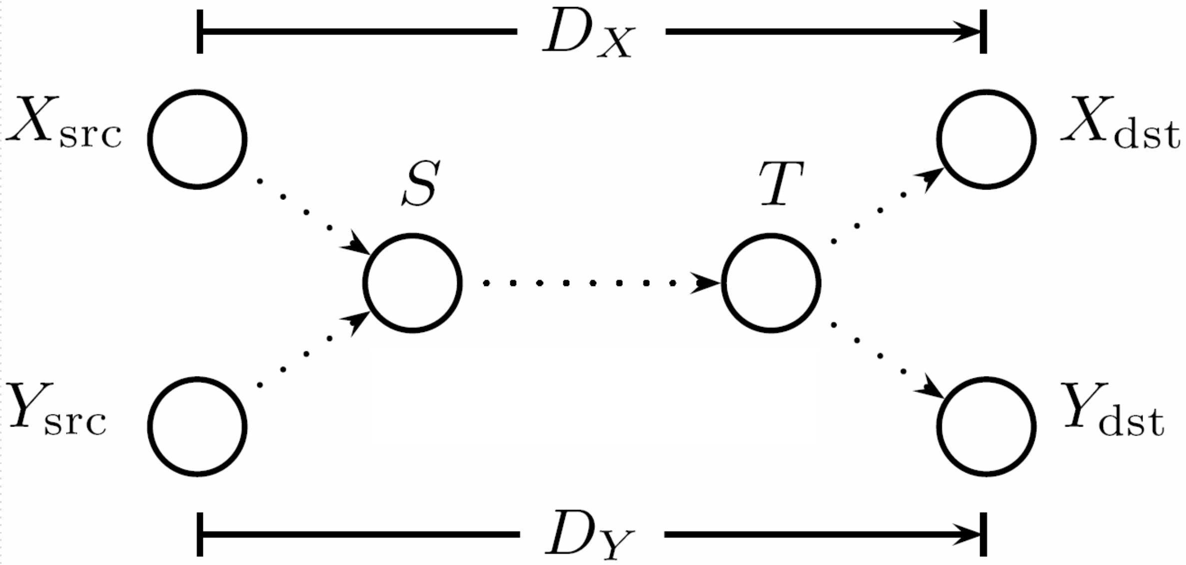

$$XCOR_{XY} = \frac{\sum_{i=1}^{n}(X_i - \overline{X})(Y_i - \overline{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \overline{X})^2 \cdot \sum_{i=1}^{n}(Y_i - \overline{Y})^2}}$$

Shared Congestion

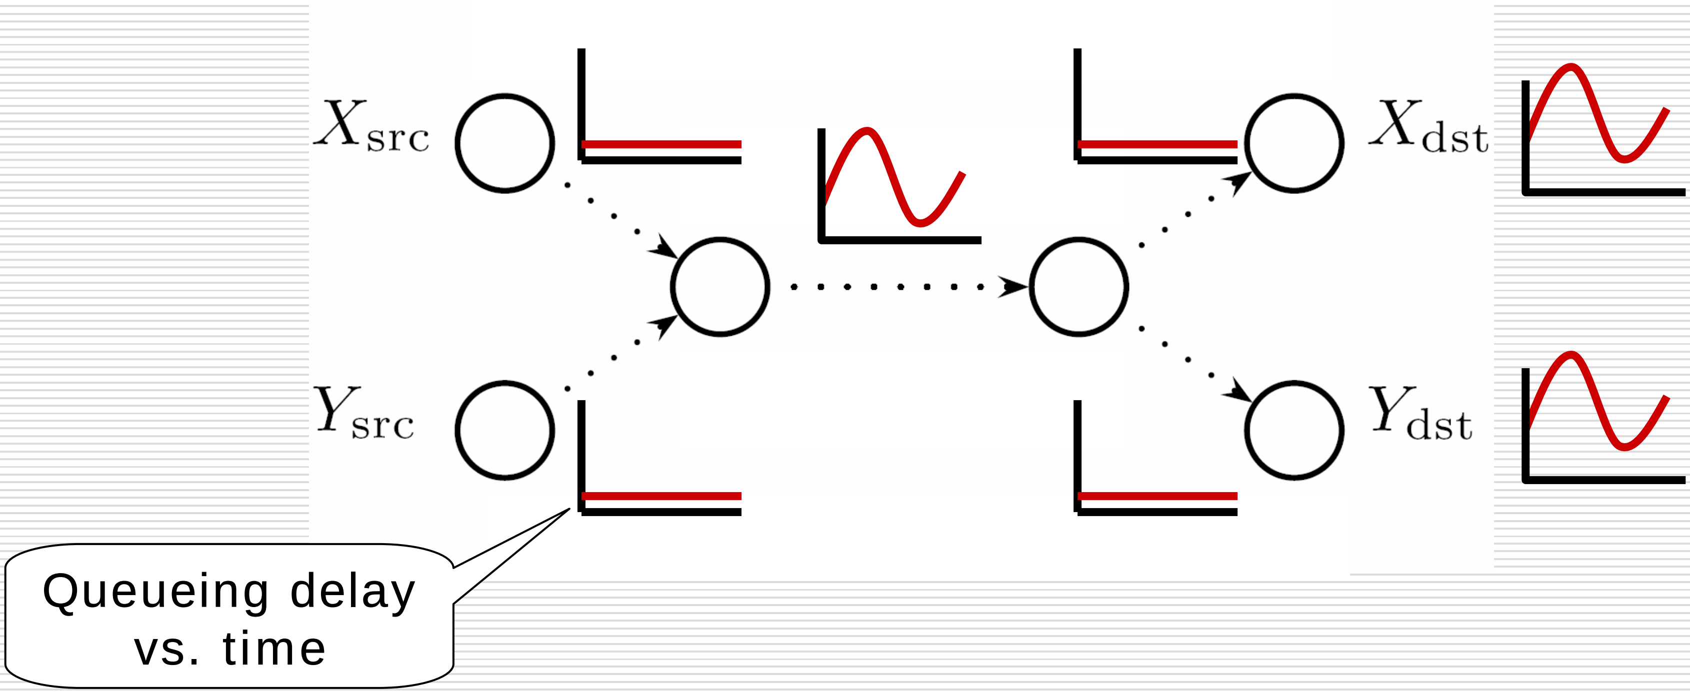

$$XCOR_{XY} = 1$$

Independent Congestion

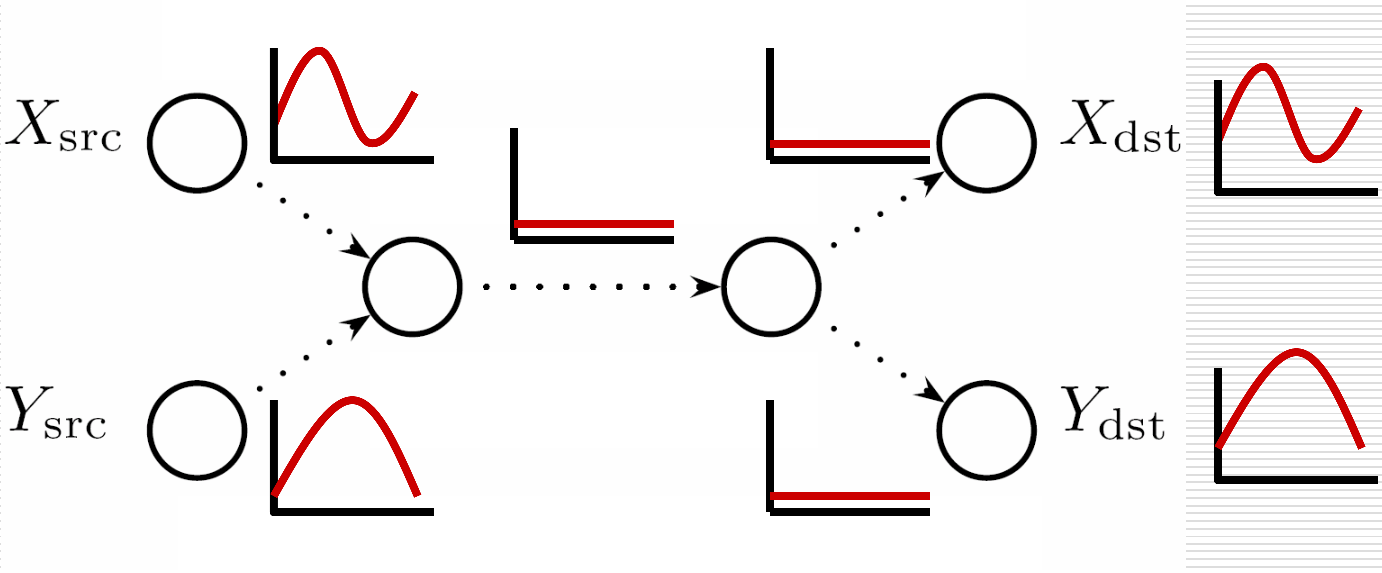

$$XCOR_{XY} \approx 0$$

1st Limitation of Basic Technique

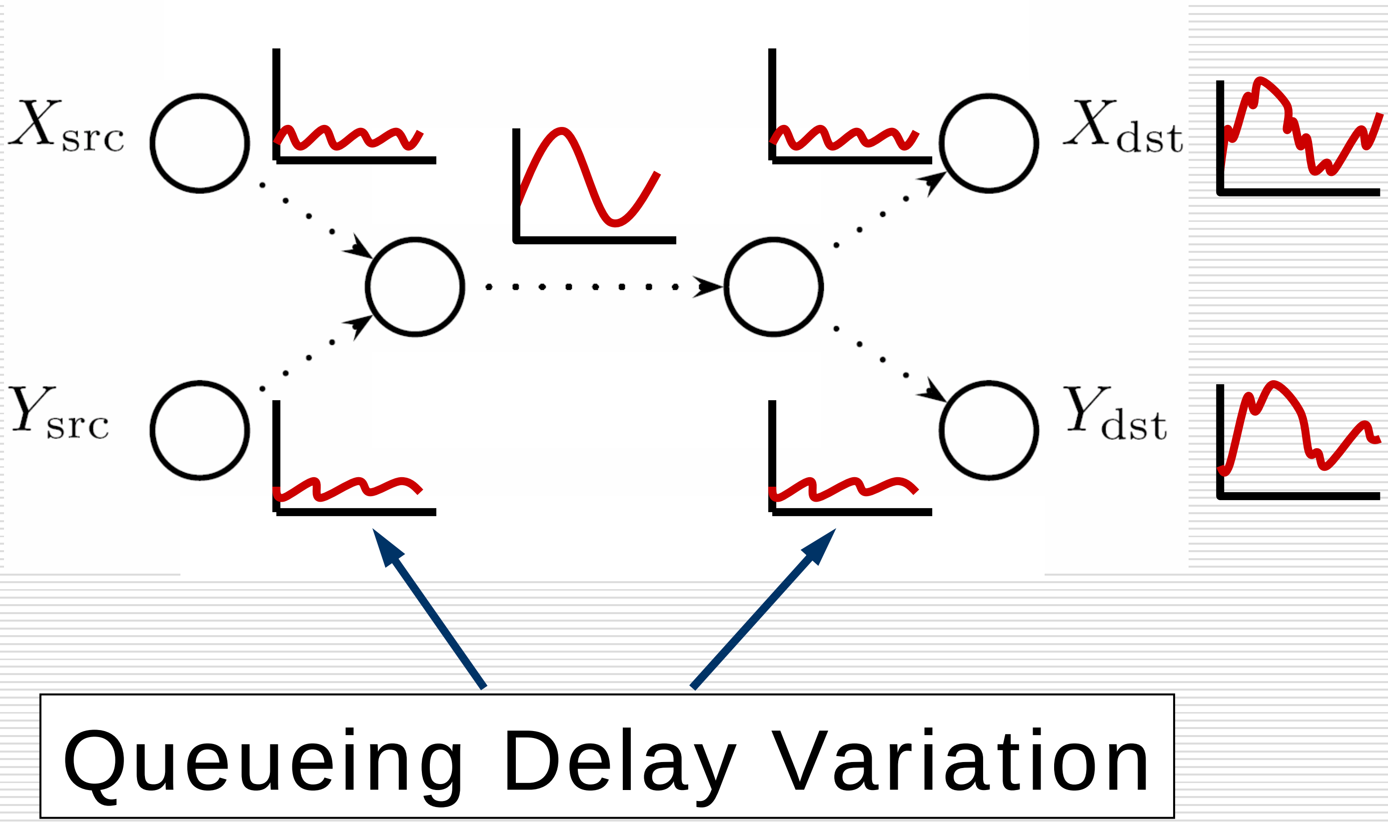

2nd Limitation of Basic Technique

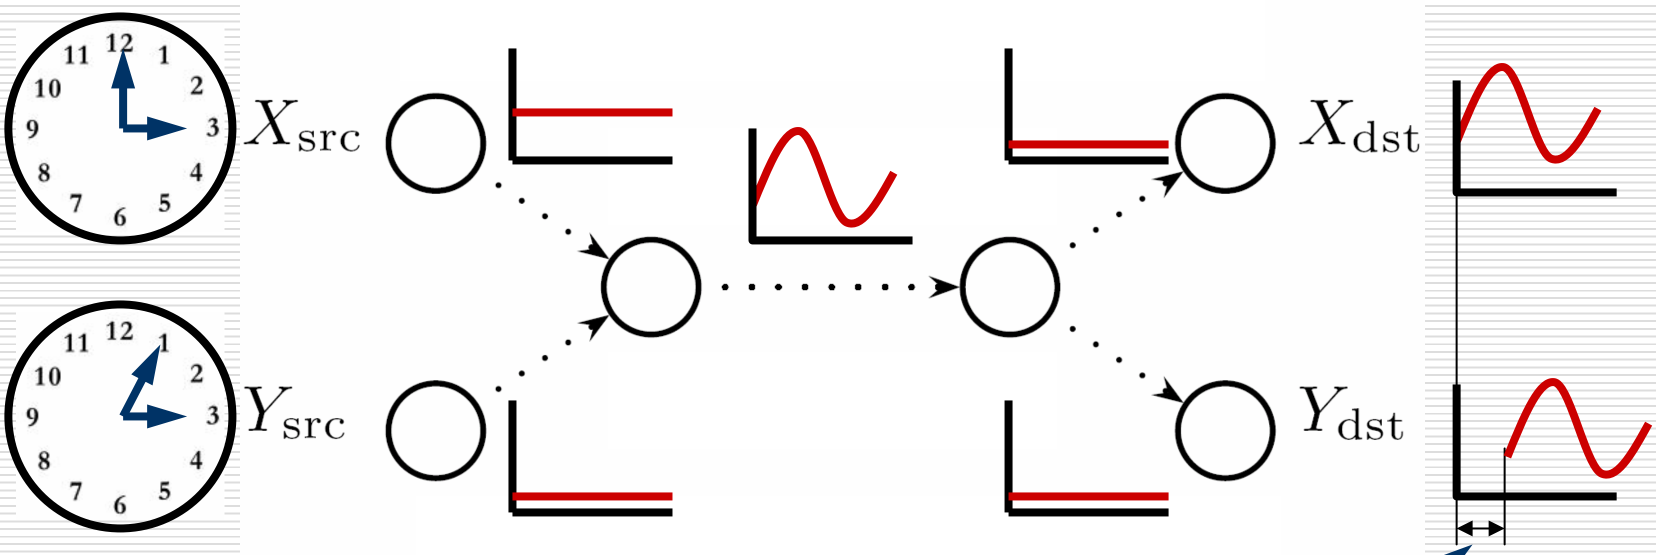

Synchronization Offset

Outline

- ☑ Introduction
- ☑ Basic technique
- ☑ Limitations of the basic technique
- ☐ DCW: Delay correlation with Wavelet denoising
- ☐ Experimental results
- ☐ Summary

Queueing Delay Characteristics

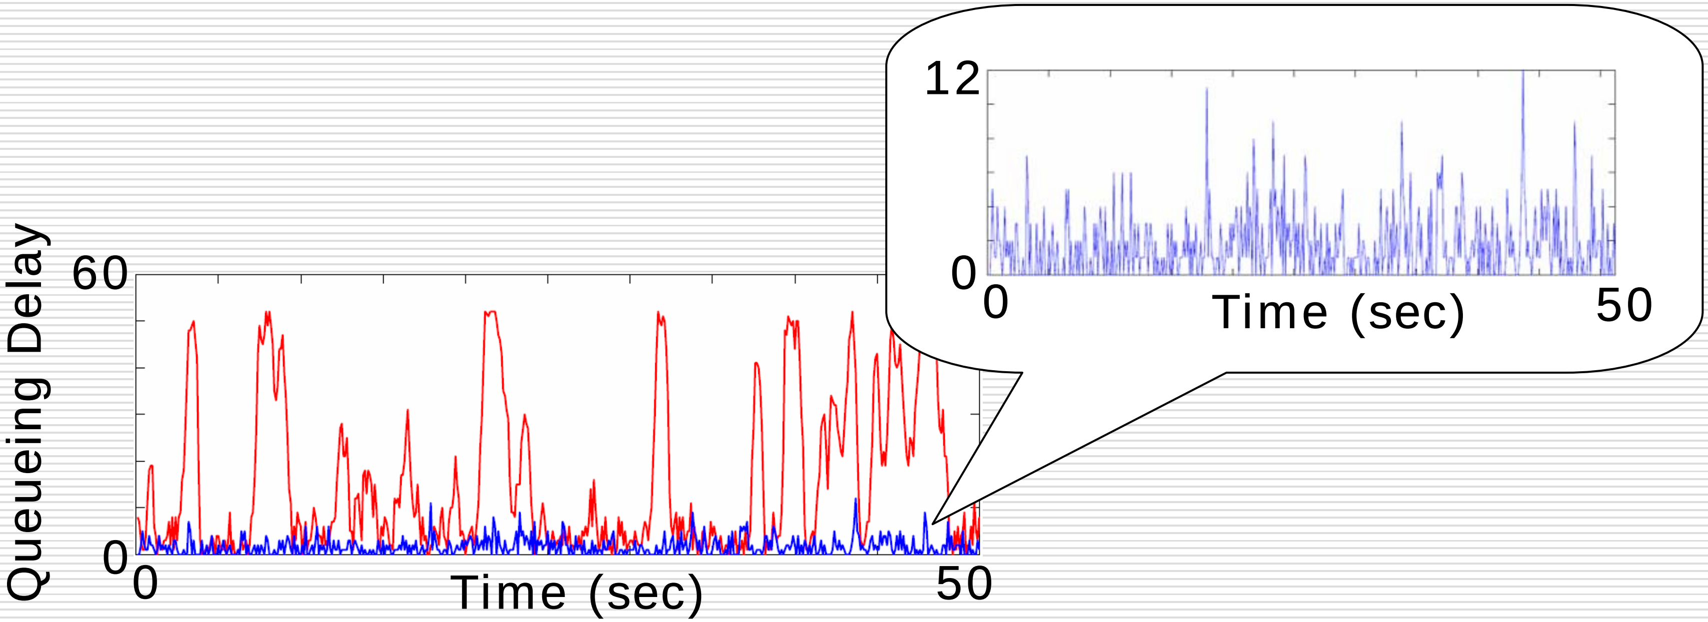

- Heavy traffic: 2%–10% loss
- Light traffic: no loss

Wavelet Transform

Measured data $x(t)$

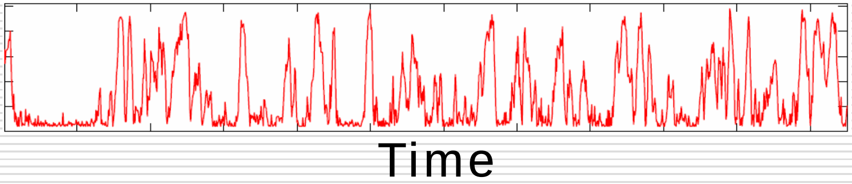

Wavelet basis $\psi_{i,j}(t) = 2^{-i/2}\psi(2^{-i}t - j)$

$$x(t) = \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} X_j^i \psi_{i,j}(t)$$

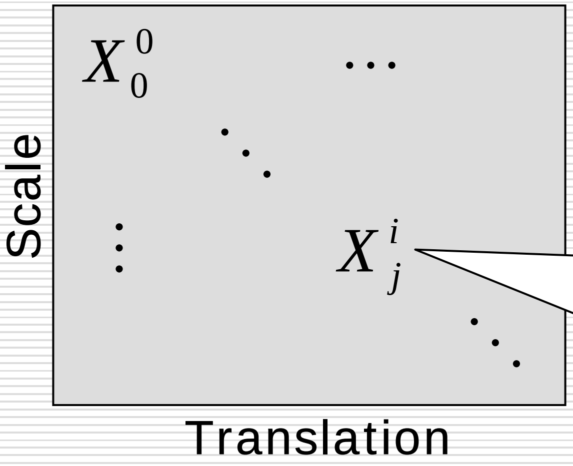

Wavelet coefficient at scale *i* and translation *j*

$$X_j^i = \int_{-\infty}^{\infty} x(t)\psi_{i,j}(t)\, dt$$

Wavelet Denoising

- Soft thresholding
  - Threshold: *T*

$$d_T(X) = \begin{cases} X - T & \text{if } X \geq T \\ X + T & \text{if } X \leq T \\ 0 & \text{if } |X| < T \end{cases}$$

$$x(t) = \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} X_j^i \psi_{i,j}(t) \Rightarrow \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} d_T(X_j^i) \psi_{i,j}(t)$$

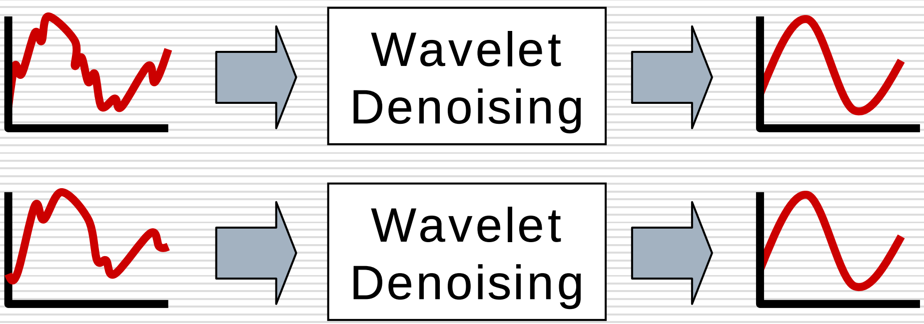

Minimizing Sync Offset Effects

- Error introduced by sync offset
  - $f(t)$: original data
  - $f(t-\Delta)$: shifted data due to sync offset
  - $f(t)-f(t-\Delta)$: error
- To minimize effects of sync offset:
  - $f(t)$ and ψ should match closely
  - $f(t)-f(t-\Delta)$ and ψ should not

Match Between Data Signal and Wavelet Basis

- Elliptic curve representation on time-frequency plane
  - C, D_1: Data Signal
  - C, D_2: Wavelet basis

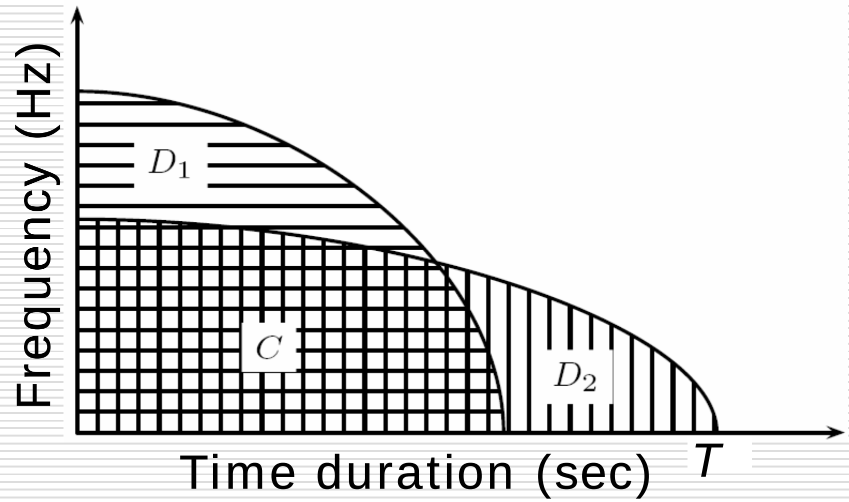

- ISNR: similarity of elliptic curves

$$\mathrm{ISNR} = \frac{1}{T} 10 \log_{10} \frac{C}{D_1 + D_2}$$

Wavelet Basis Selection

- Differential ISNR
  - (ISNR between $f(t)$ and ψ) – (ISNR between $f(t)$-$f(t$-$\Delta)$ and ψ)
- Daubechies wavelets
  - Simple
  - Easy to implement

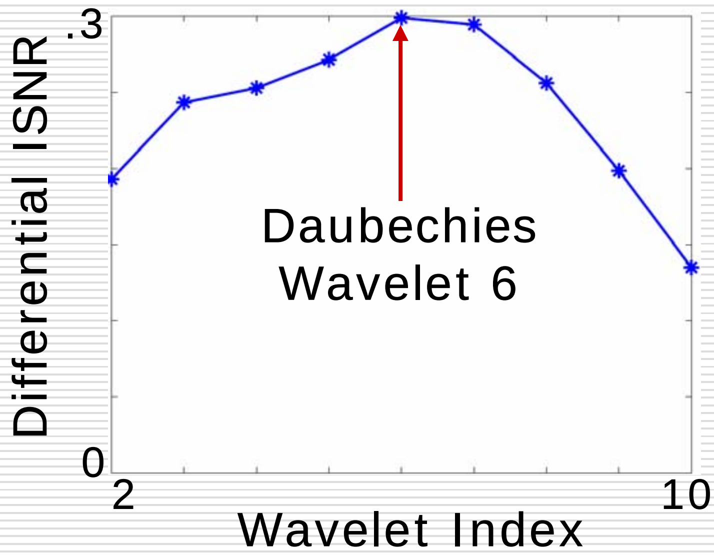

Evaluation

- Comparison with
  - MP: delay-based [Rubenstein, et al]
  - BP: loss-based [Harfoush, et al]
- Positive Ratio

$$\frac{\text{\# of answers indicating shared congestion}}{\text{\# of experiments}}$$

  - 1: shared congestion
  - 0: no shared congestion

Common Source Topology

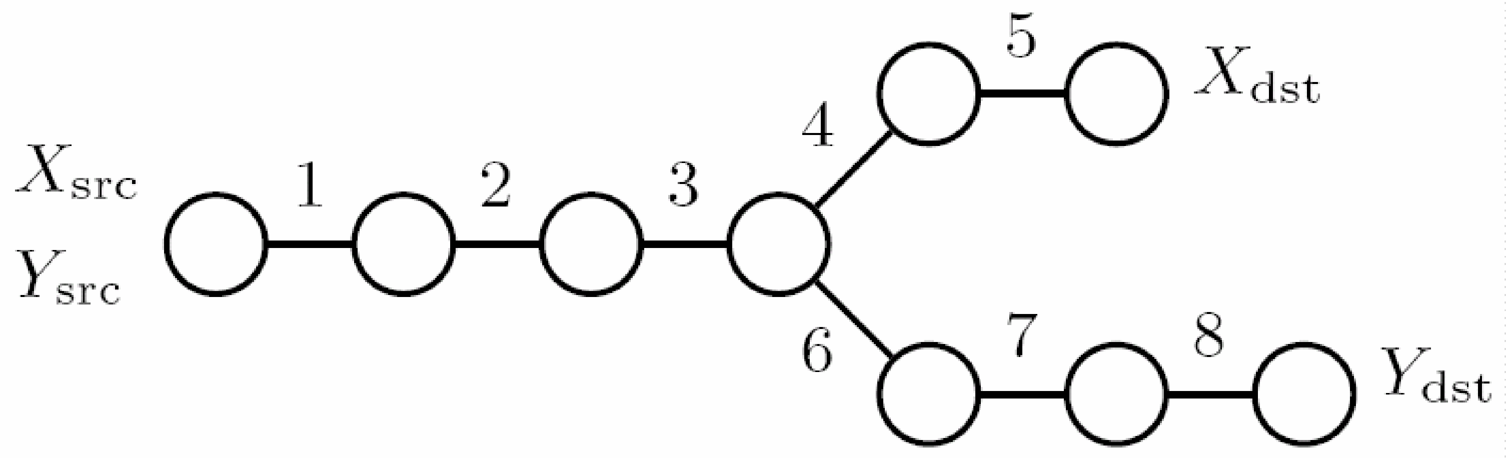

- X_{src} and Y_{src} are synchronized
- No synchronization offset

Common Source / Drop-Tail / Long-Lived TCP Traffic

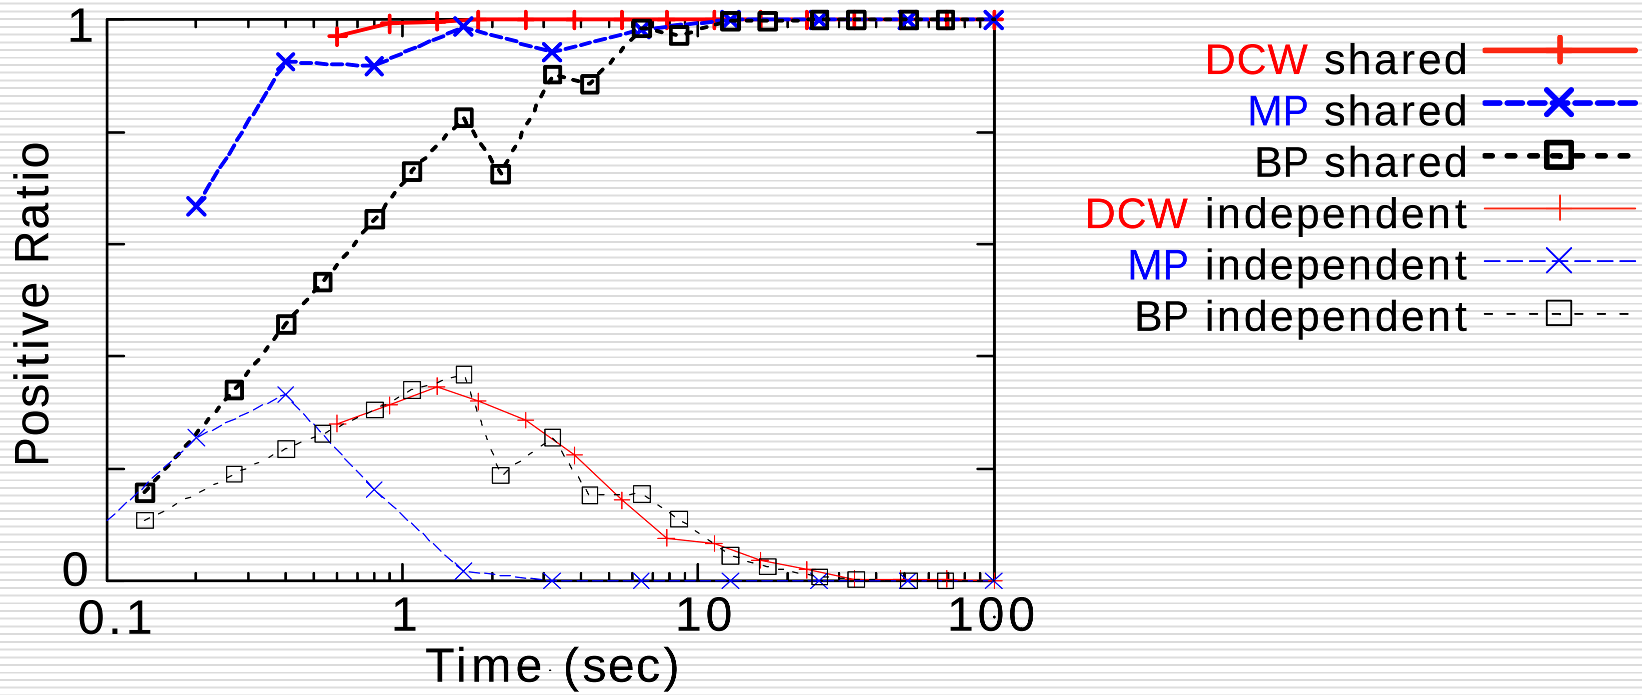

- Shared: DCW $\succ$ MP $\succ$ BP
- Independent: MP $\succ$ DCW $\approx$ BP

Common Source / Drop-Tail / On-Off CBR Traffic

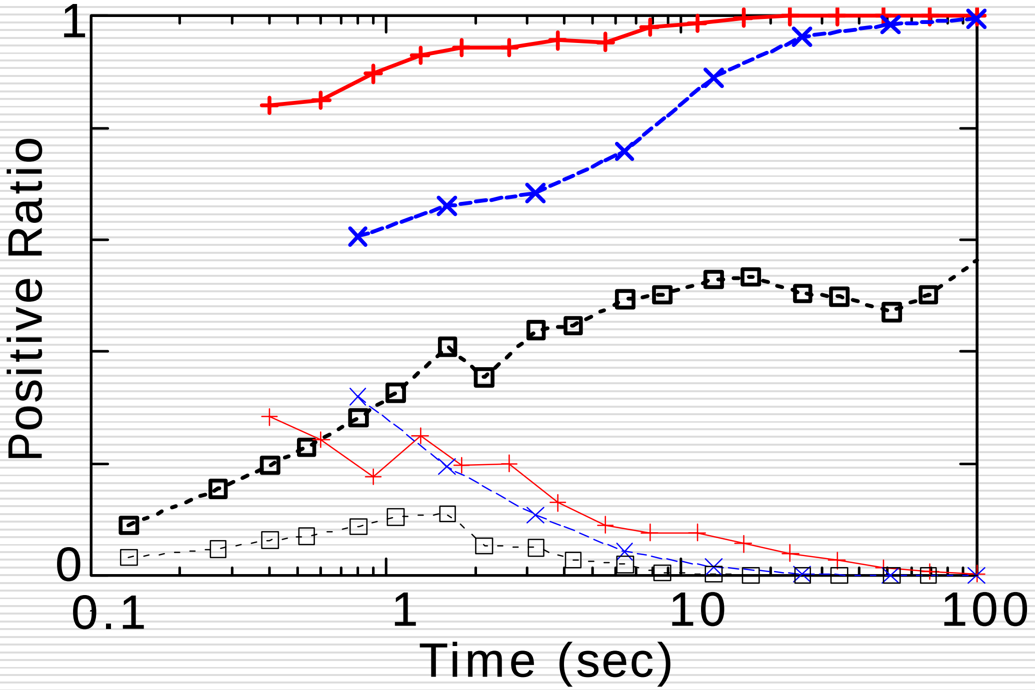

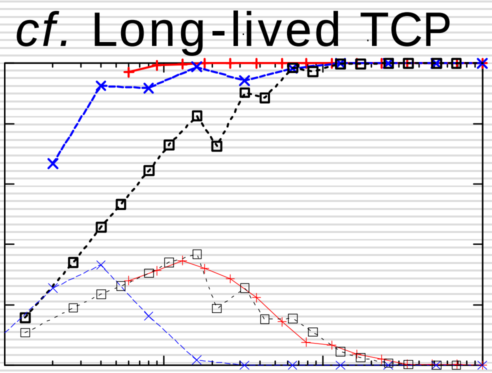

- Slower convergence due to
  - Delay on non-congested links → DCW, MP
  - Shorter loss runs → BP

Common Source / Drop-Tail / Short-Lived TCP Traffic

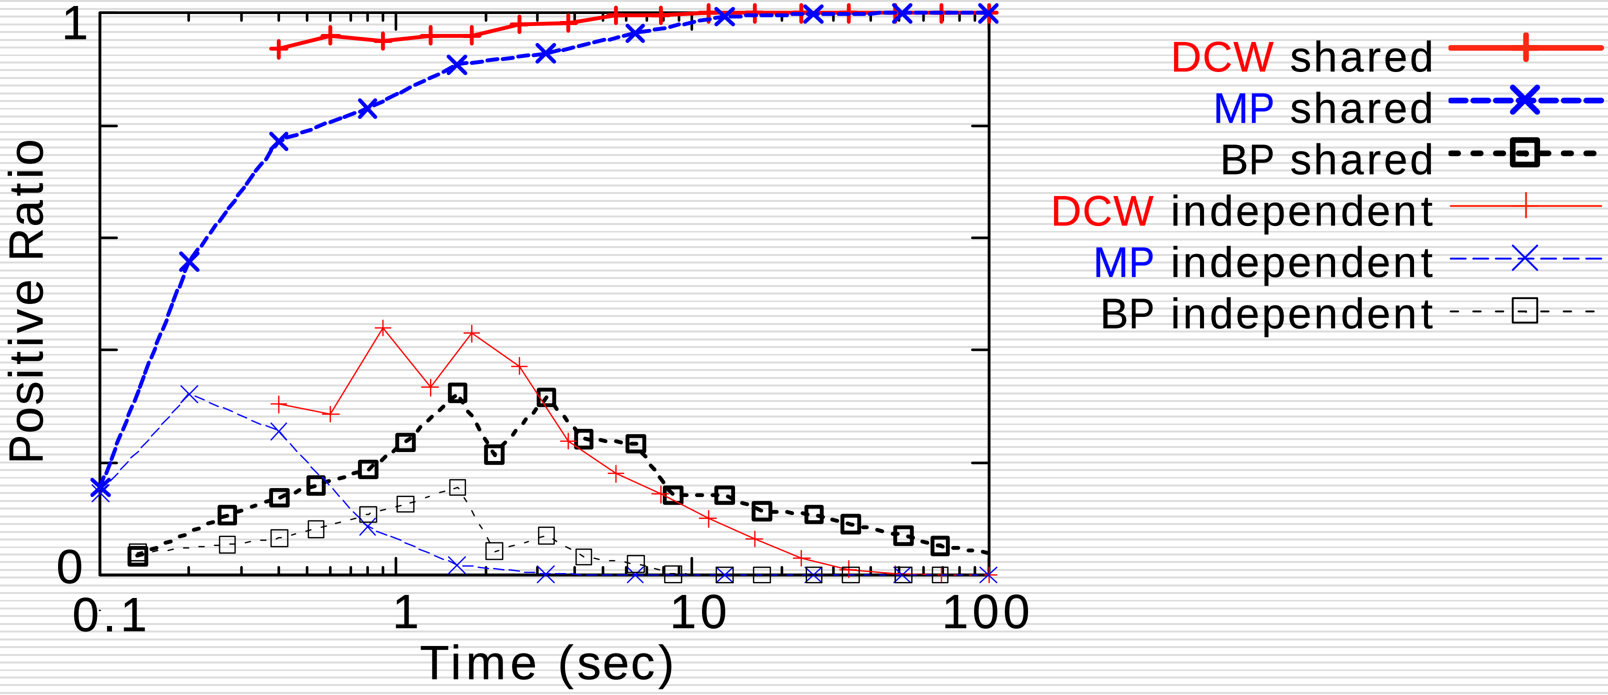

- Even shorter loss runs → BP fails.

Common Source / RED

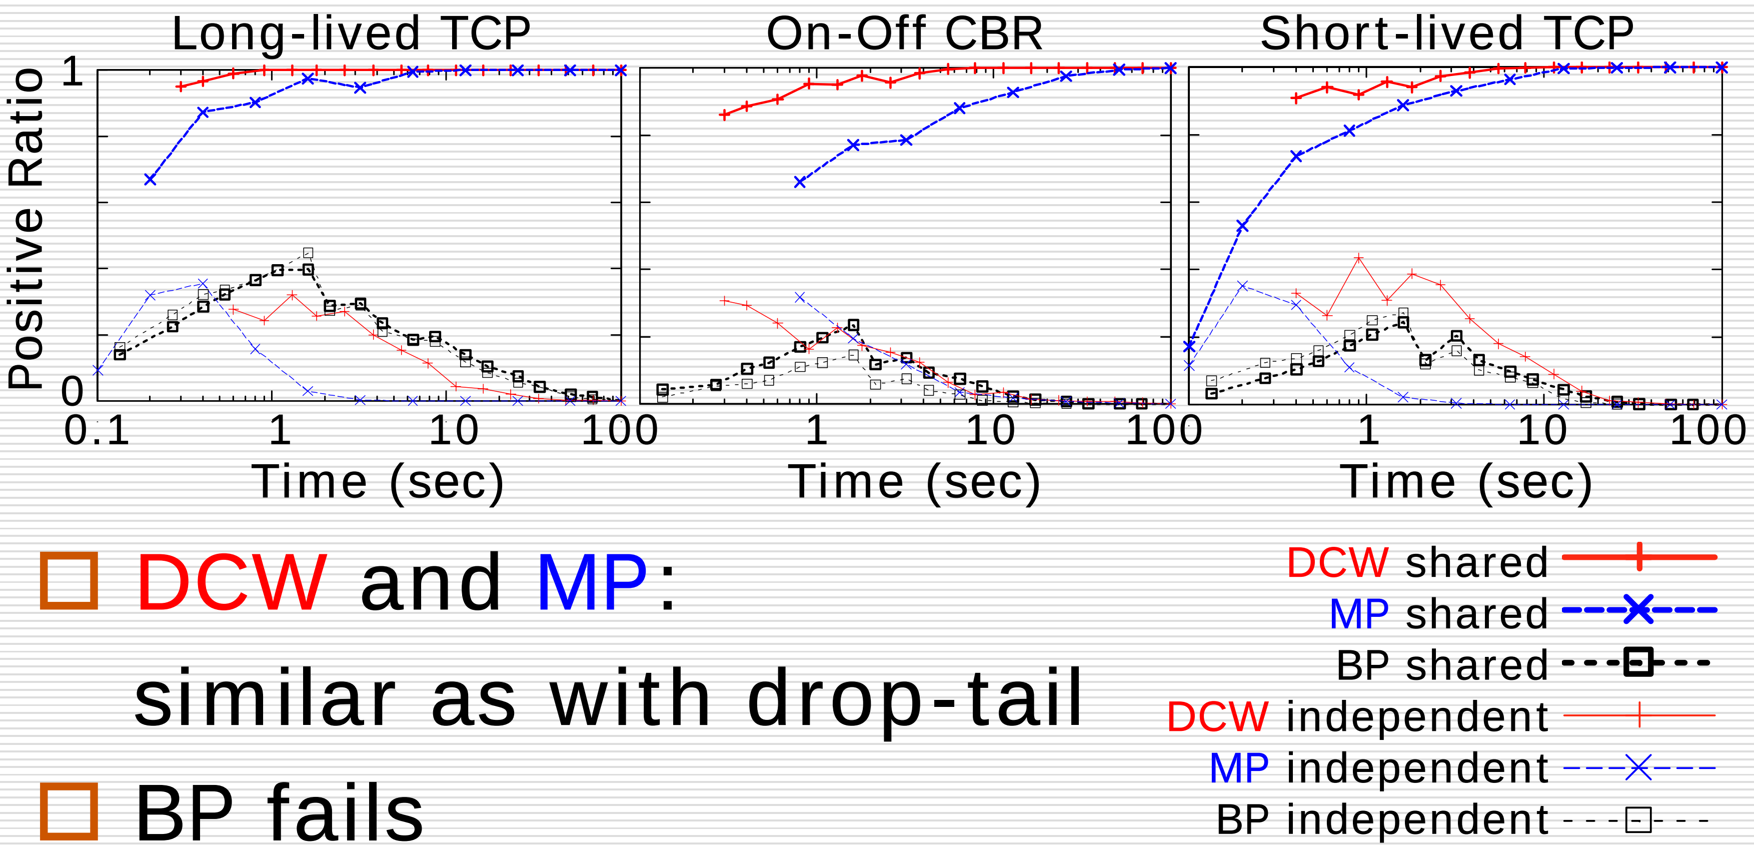

- DCW and MP:
  similar as with drop-tail
- BP fails

Topology without Sync Point

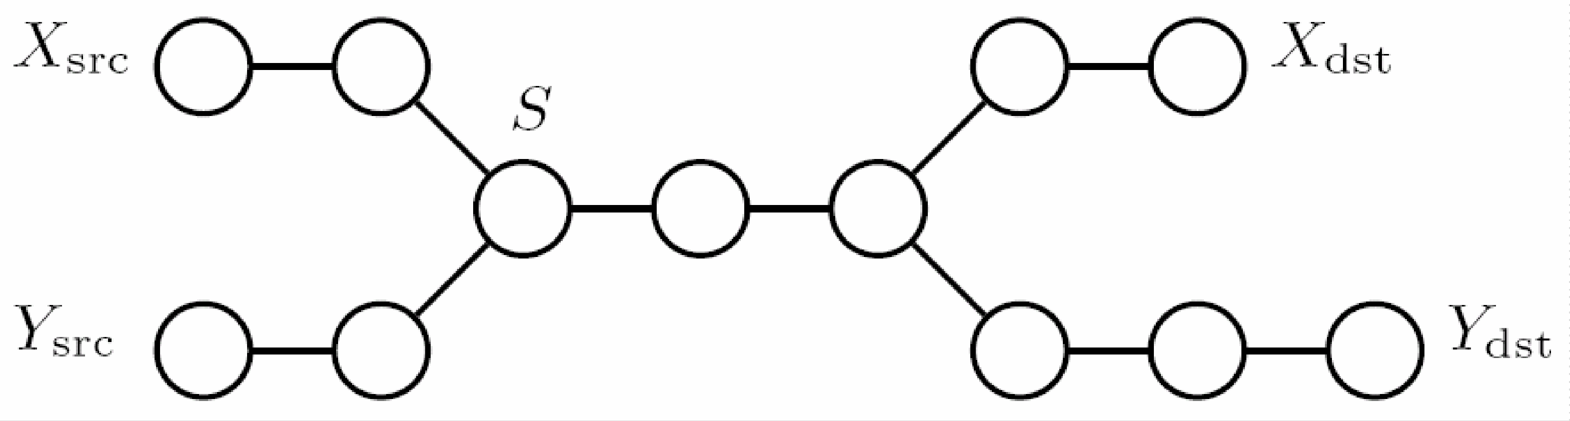

- Synchronization offset > 0

Sync Offset Tolerance

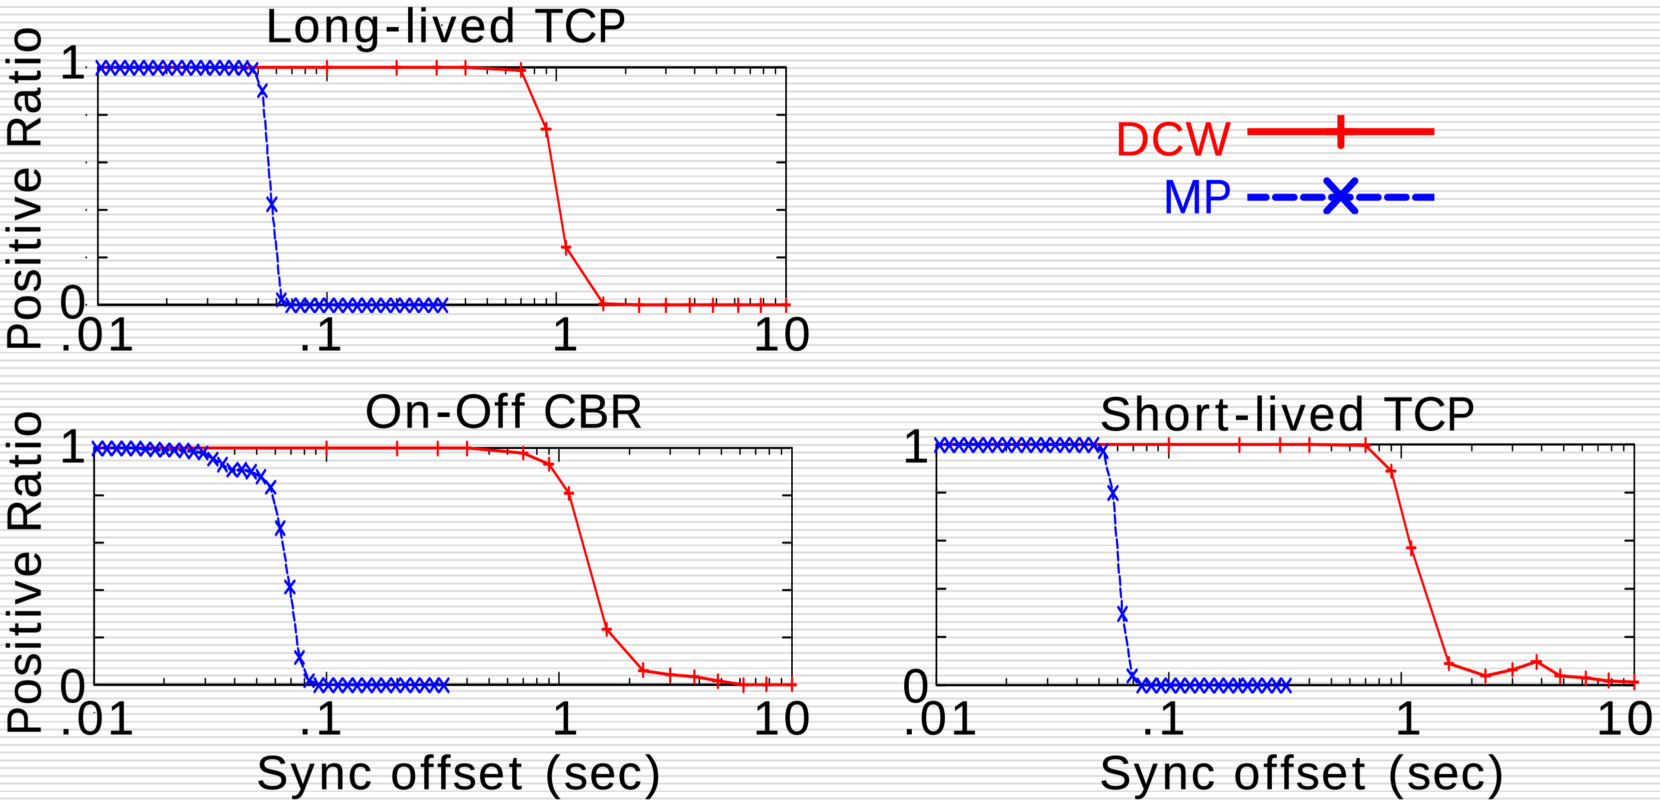

- DCW: 1–2 sec, MP: 30–70ms, BP: < 10ms

Internet Experiment

- Topology

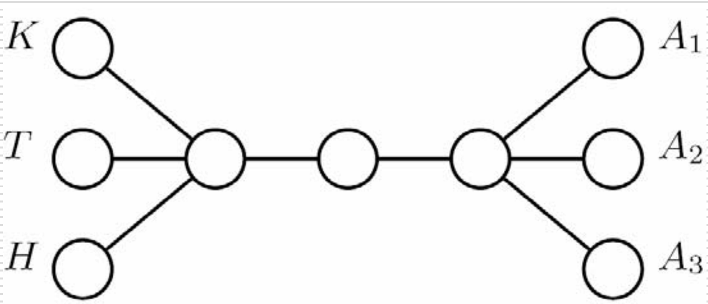

- 10 seconds to converge

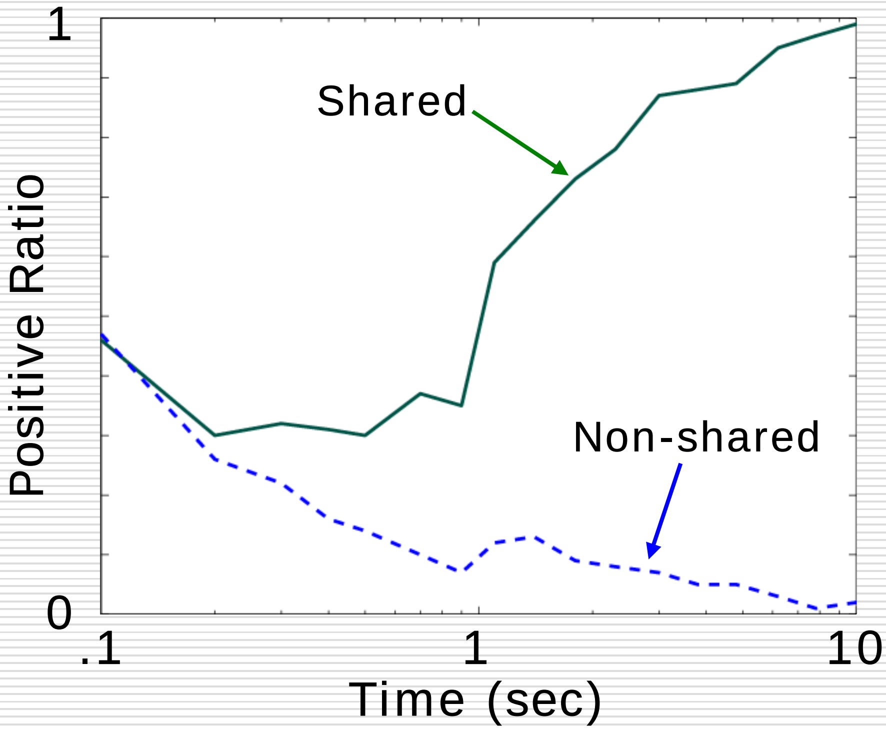

Summary

- Proposed technique: DCW
  - Delay Correlation with Wavelet denoising
- As fast and accurate as previous techniques (with a common endpoint)
- Applicable to *any 2 Internet paths*
- Basic primitive for overlay topology improvement