

Empirical Bulk Transfer Capacity

Matt Mathis

mathis@psc.edu

<http://www.psc.edu/~mathis>

Mark Allman

mallman@lerc.nasa.gov

<http://gigahertz.lerc.nasa.gov/~mallman>

IPPM IETF-43

3-Dec-1998

Empirical Bulk Transfer Capacity

BTC is a measure of a network's ability to transfer significant quantities of data with a single congestion-aware transport connection (e.g., TCP).

"Standard" Congestion control is crucial

... but the present is a moving target

RFC2001.bis (Allman et al)

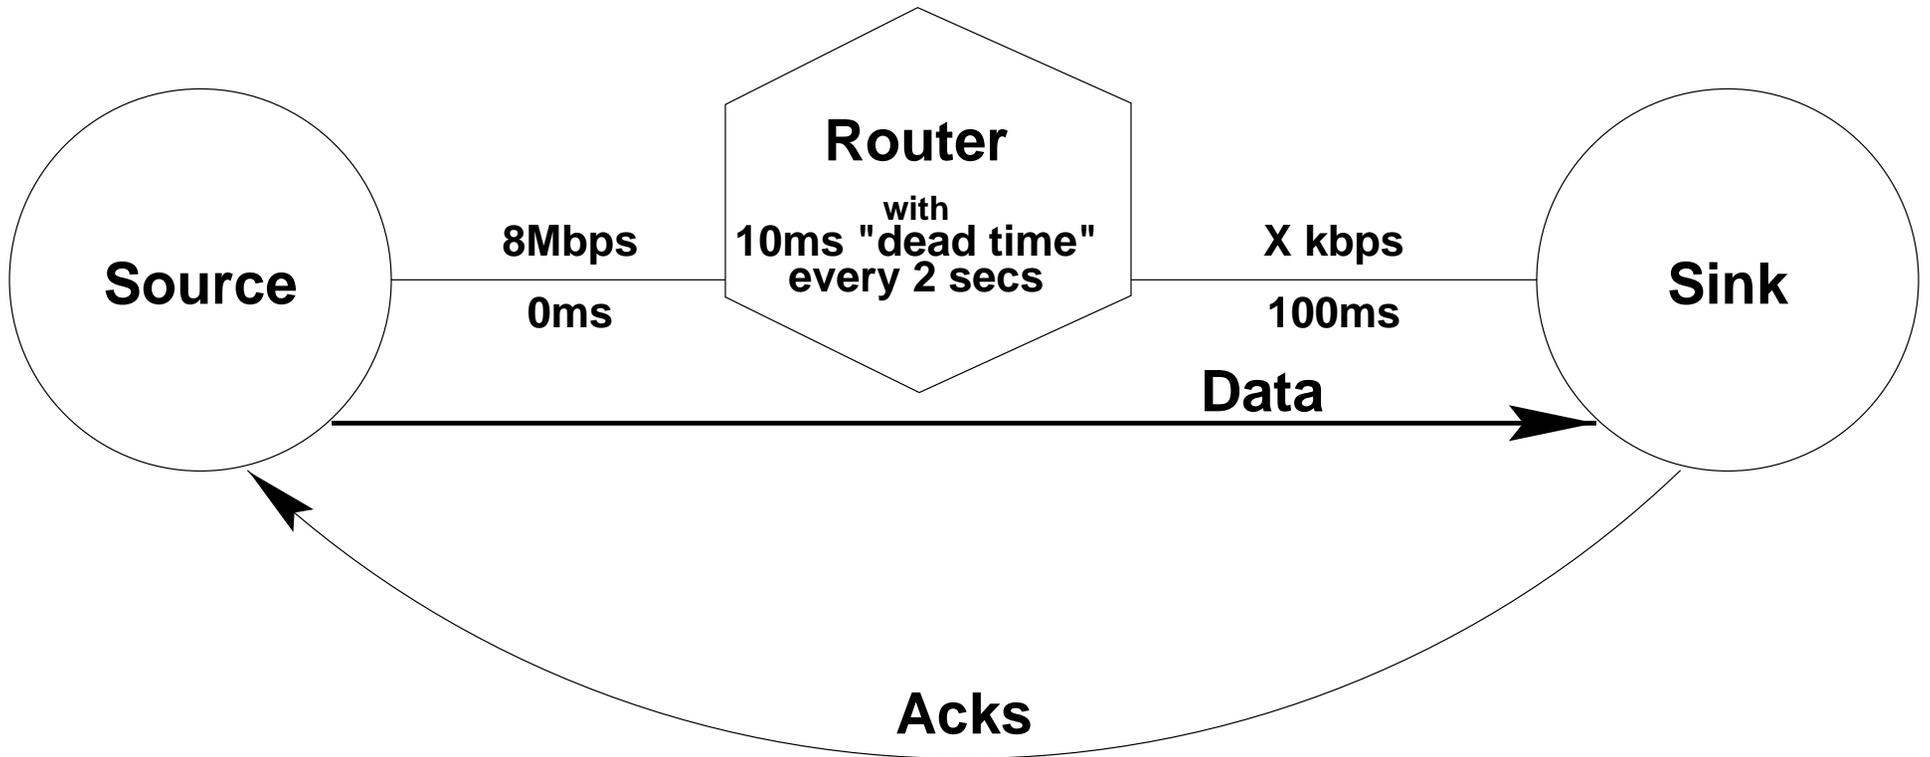
NewReno (Floyd & Henderson)

ECN (Ramakrishnsnan & Floyd)

A problematic example

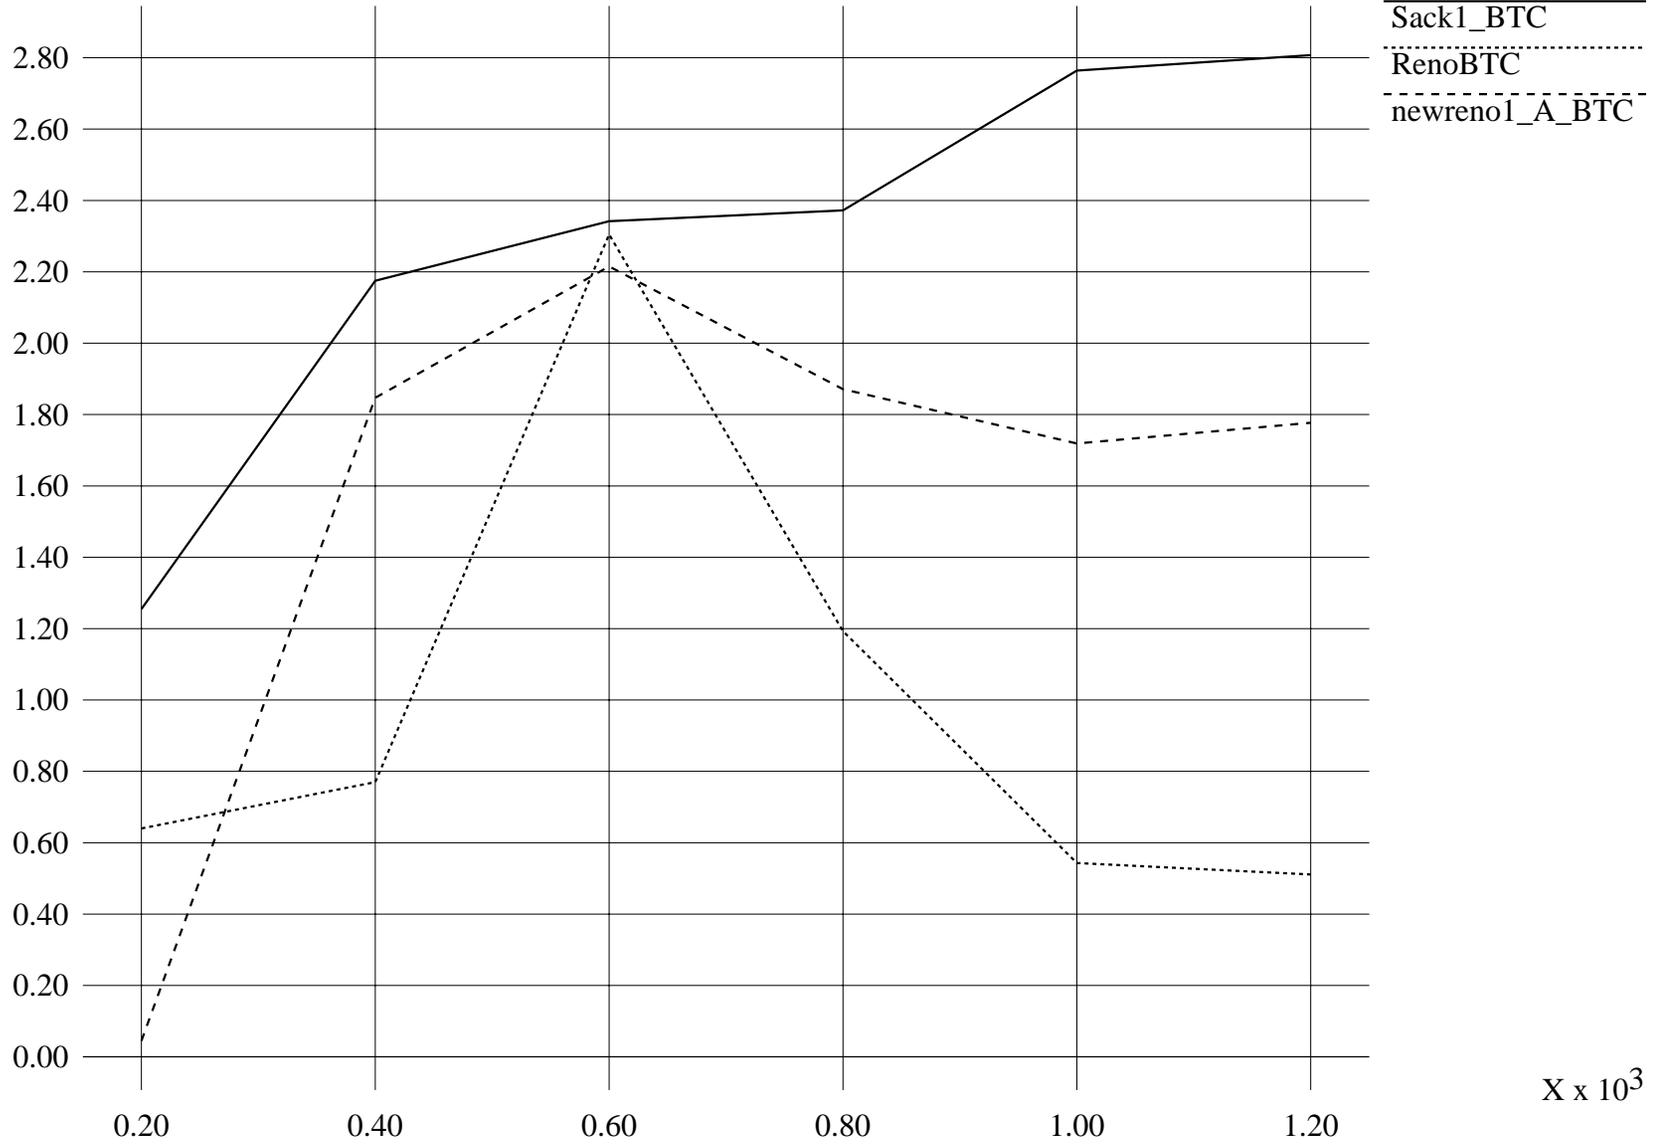
Reno TCP fails as a metric because "better" paths sometimes yield "poorer" measures due to fast recovery failures

A non-monotonic metric does not make sense



X Graph

Y x 10³



X x 10³

Three pronged approach

- Allow multiple BTC metrics
- Tight specifications for BTC metrics
(Reference implementations?)
- Include additional ancillary metrics

Ancillary Metrics

Use models to corroborate the BTC

[Mathis, CCR July 1997]

[Padhye, SigCOMM 1998]

?? Transport only path properties

(e.g. queue size, packet reordering, burst losses)

Path pathologies

(e.g duplicate packets, load induced outages)

Revisit metric requirements

"Better" paths had better yield "better" measures

Need A-frame to unify simple metrics (loss delay, etc)
with BTC and ancillary metrics

Consider 2 different BTC metrics
with nonuniform sensitivity to reordering

Can each be estimated from:

- the other plus a reorder metric

or from

- delay, loss and reorder metrics?

Frees us from requiring the answer to the research question
"how much reordering is ok"