SIEMENS

Comparison of TCP Performance Models

[July 23, 01]

Hans-Peter Schwefel / Manfred Jobmann / Daniel Höllisch / Dan Heyman

Siemens ICM / Technische Universität München / AT&T Labs

Hans.Schwefel@icn.siemens.de

Motivation: Internet on Air (IoA)

IoA: The Evolution Target of Mobile Networks and Internet

Trends in Mobile Networks:

- IP transport in the backbone; IP in RAN
- Transport voice & data over IP
- Terminate IP in the mobile host
- Separation: Transport ←→ Control

Trends in the Internet:

- Enable wireless access, support mobility
- QoS beyond “Best Effort”
- Security and AAA

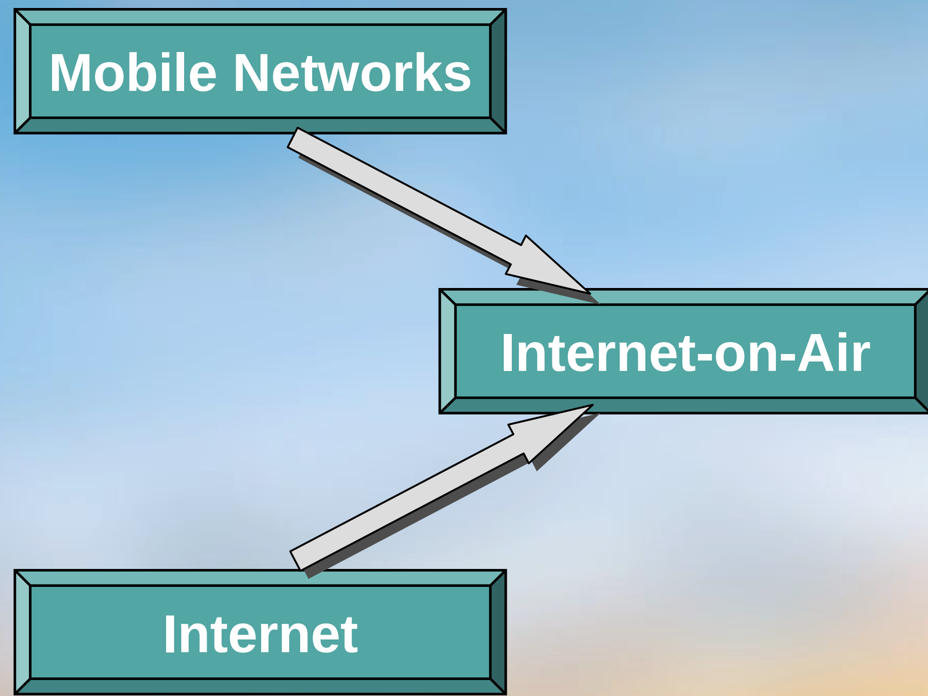

TCP traffic dominant in the Internet → Wireless TCP important in IoA

Challenge: Traffic/Performance Models for TCP

[Even without wireless links]

End-to-End flow-control (using Ack packets)

→ Feedback: network behavior (congestion) ↔ ingress traffic

→ No separation traffic model/network model possible

→ New traffic/performance models required

Here: Comparison of 3 models & simulation (TCP Reno)

- TCP-UMass [Padhye et al., Sigcomm 98]
- TCP-Engset [Heyman et al., Sigmetrics 97]
- TCP-NBurst [modific. of Schwefel, Infocom 01]

Network Scenario I: Multiplexed ON/OFF Traffic [without wireless link]

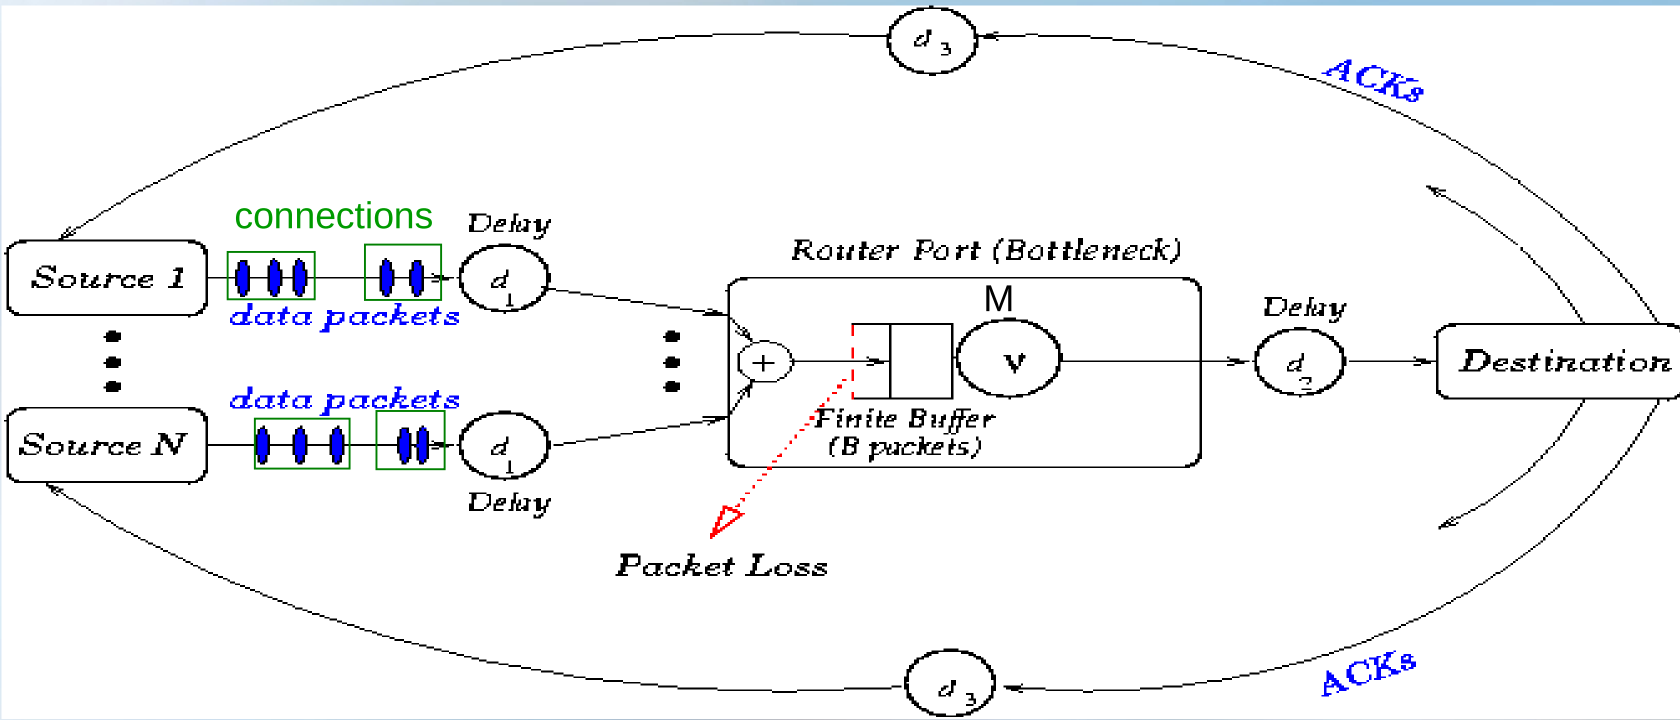

- Scenario: N (LRD) ON/OFF sources, queueing/loss only at bottleneck router
- Models without flow-control → known results in literature
- Here: Investigate average throughput per connection in TCP setting
- Two scenarios: (A) fast access, slow trunk, (B) vice versa

Network Scenario II: Parameter Settings

	Scenario (A)	Scenario (B)
access link speed	10 Mbit/s	128 kbit/s
trunk speed	1 Mbit/s	1.5 Mbit/s
packet size (fixed)	1000 Byte	576 Byte
average conn. size	50 kByte	200 kByte
access link speed	$\lambda_p = 1250$ packets/s	$\lambda_p = 27.78$ packets/s
trunk speed	$\nu = 125$ packets/s	$\nu = 325.52$ packets/s
average connection size	$n_p = 50$ packets	$n_p = 347.2$ packets
Buffer-Size	$B = 16$ packets	$B = 50$ packets
Delays	$\sum_i d_i = 80$ ms	$\sum_i d_i = 300$ ms
resulting minimal RTT	$R_0 = 88.8$ ms	$R_0 = 339$ ms

Observed:

- **Goodput per connection *i*:** $G_i := C_i / D_i$
- **Average Goodput & weighted average:** $\overline{G} := \frac{1}{K} \sum_{i=1}^{K} G_i, \quad \widehat{G} := \frac{1}{\sum_{i=1}^{K} C_i} \sum_{i=1}^{K} C_i \cdot G_i$
- **where** C_i:=#packets in conn. *i*, D_i:=duration of conn. *i*

Analytic Models I: TCP-UMass [Padhye, Firoiu, Towsley, Kurose]

Approximation of throughput T in persistent TCP connection

- Derived from sender/receiver behavior
- Input Parameters
  - b: # packets acknowledged by ACK [here: b=1]
  - $\overline{RTT}$: Average Round-Trip-Time (including queuing delay)
  - W_{max}: Maximum size of congestion window [here: W_{max}=1024]
  - T_0: Average time-Out interval
  - p: Fraction of retransmitted packets

$$T(p,\overline{RTT}) \approx \frac{1}{\overline{RTT}} \min\left(W_{max}, \frac{1}{\sqrt{\frac{2}{3}bp} + T_0/\overline{RTT}\,\min\left(1, 3\sqrt{\frac{3}{8}bp}\right) p(1+32p^2)}\right)$$

Analytic Models II: TCP-Engset [Heyman, Lakshman, Neidhardt]

Modified Processor Sharing model on connection level:

- **For j active users, each obtains service at rate**

  ~λ_p when $j\lambda_p \leq \nu$

  ~ $\delta\nu/j$ when $j\lambda_p > \nu$

- **with attenuation factor $\delta \leq 1$ derived from TCP behavior**

$$\delta = 1 - \frac{[(1-p_{low})^+]^2}{2(p_{high}-p_{low})}, \quad p_{high} = \frac{B}{\nu R_0} + 1, \quad p_{low} = \frac{p_{high}}{2^{1.5}}$$

With R_0 = minimal Round-Trip Time = $\Sigma d_i + 1/\nu + 1/\lambda_p$

Note:

- **# active connections can be represented by birth-death process**
- **Steady-state prob. independent of distribution of connection size (→ no impact of Long-Range Dependence by PT conn. sizes)**

Analytic Models III: TCP-NBurst [Schwefel]

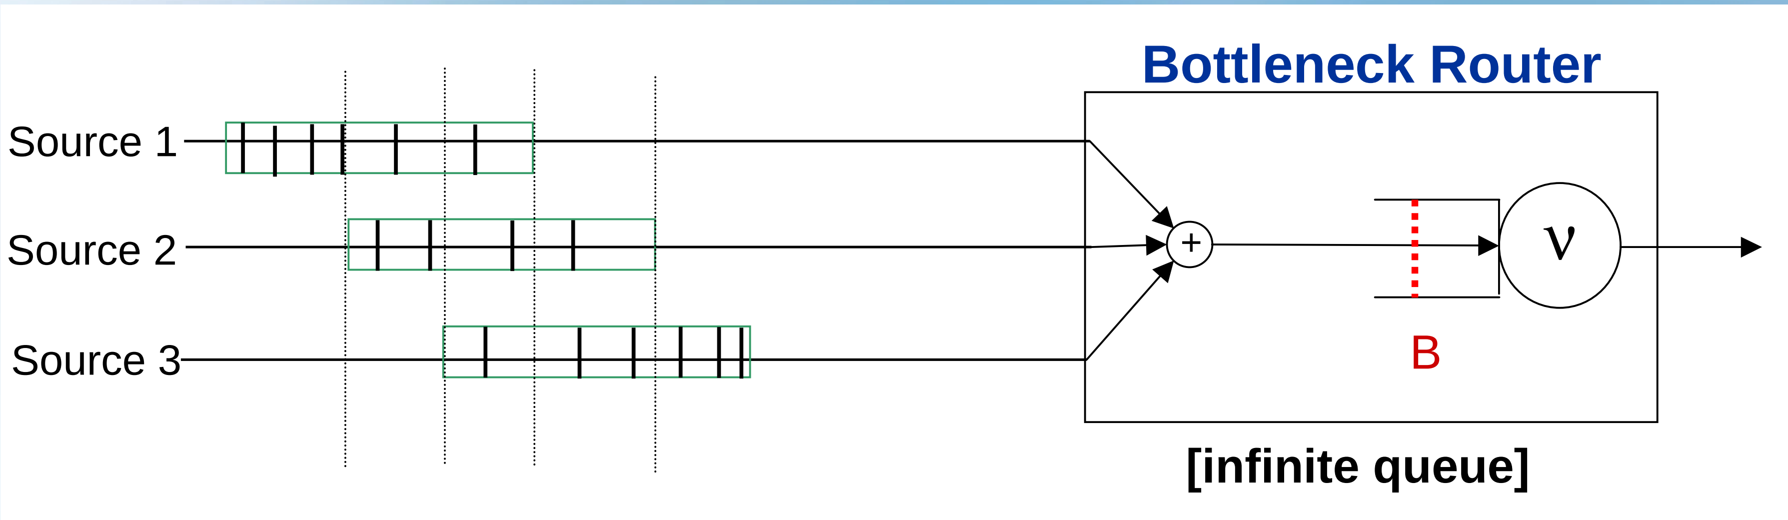

Cmp. TCP-Engset

Packet-level extension of TCP-Engset model:

- ‚Sharing' of bandwidth: packet-rate λ_p at source reduced to $\delta\nu/j$ for j active sources when $j\lambda_p>\nu$
- conn. duration extended → #packets in conn. unchanged
- Throttling only during congestion state when Q>=B

Comparison (I): Applicability of Models

- **TCP-UMass**

  + Simple Formula, based on end-to-end TCP behavior

  -- Computation of throughput only

  -- ‚Derived' parameters required: $\overline{RTT}$, T_0, p

  ? assumes persistent connections

- **TCP-Engset**

  + use of basic parameters only

  + dynamic connections (ON/OFF)

  -- no packet-level queue → restricted performance parameters (e.g. loss rate)

  ? Independence of connection size distribution (LRD properties)

- **TCP-NBurst**

  + packet-level queue → wide range of computable performance parameters

  -- computationally hard (matrix-algebraic methods), potential numerical problems

Comparison (II): Goodput per connection [Scenario (A)]

Geometric Connection Size

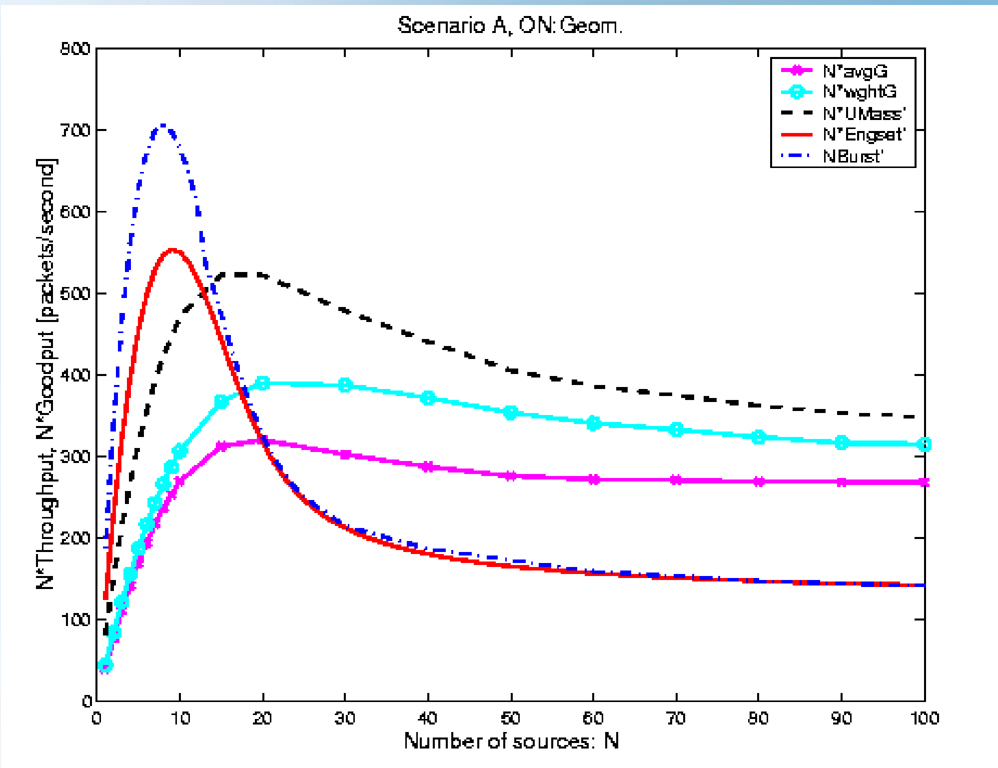

Power-Tailed Connection Size (LRD)

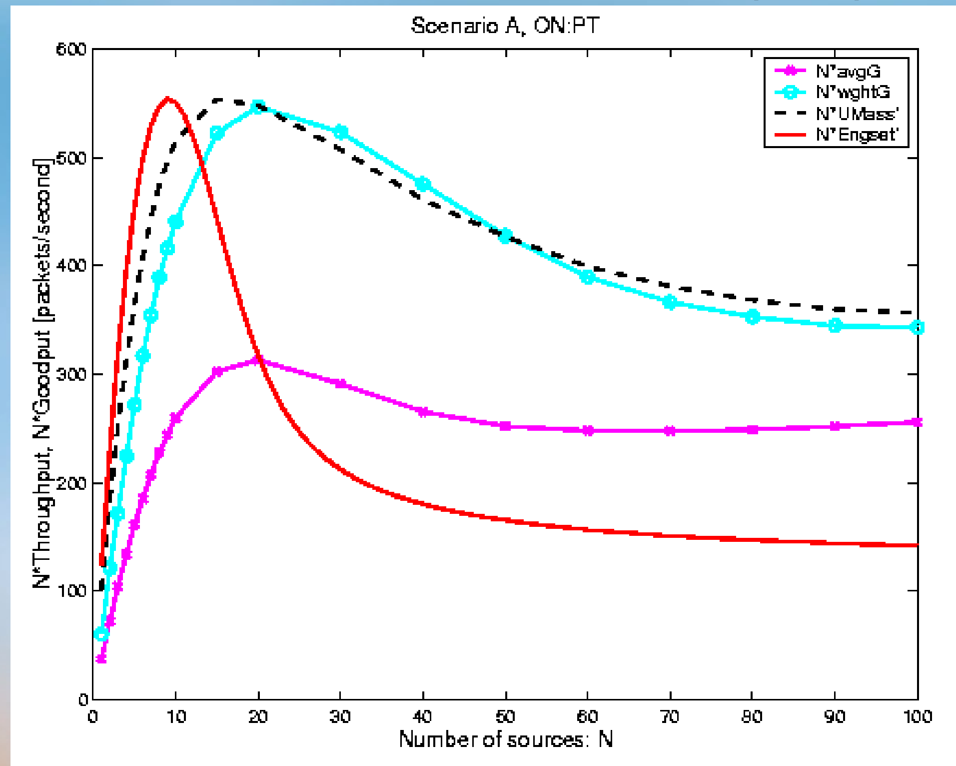

- **TCP-UMass closest fit**
- **TCP-Engset and TCP-NBurst equivalent for large N (but underestimation)**

- **Weighted Goodput >> Plain Average**
- **Excellent fit of TCP-UMass to weighted average**
- **TCP-NBurst not computable here**

SIEMENS

Comparison (III): Goodput per conn. [Scen. (A), N=20]

Variation of Buffer-Size

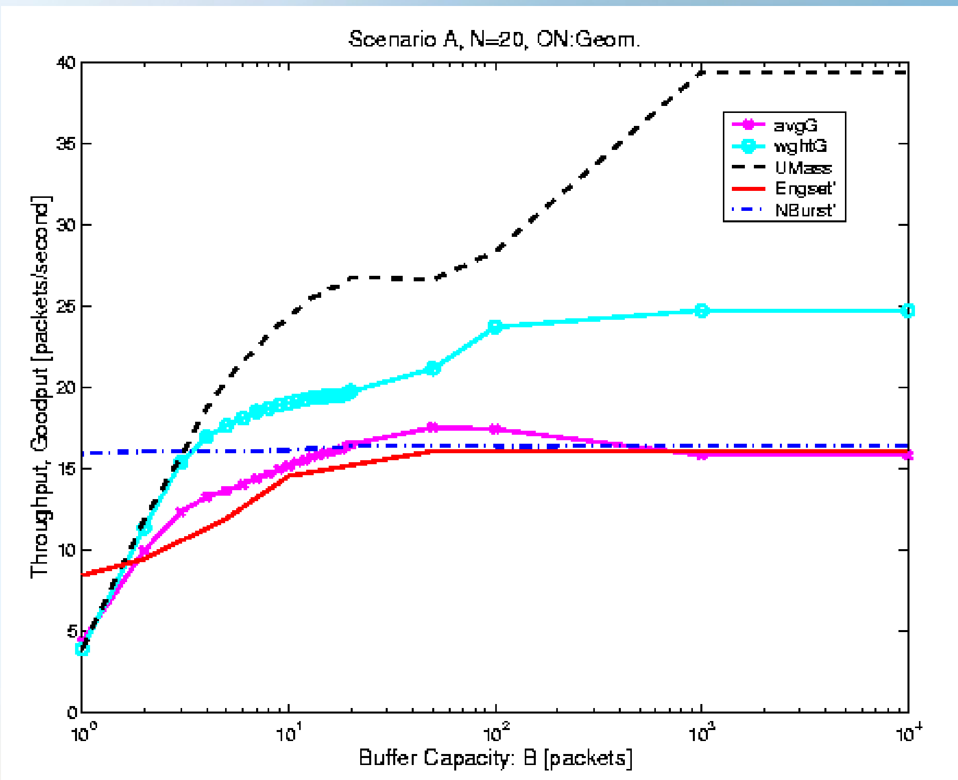

- **Small buffers B<5: TCP-UMass good**
- **Large buffers: TCP-Engset and TCP-NBurst match plain average**
- **TCP-Engset best match for shape of curve**

Variation of av. connection size

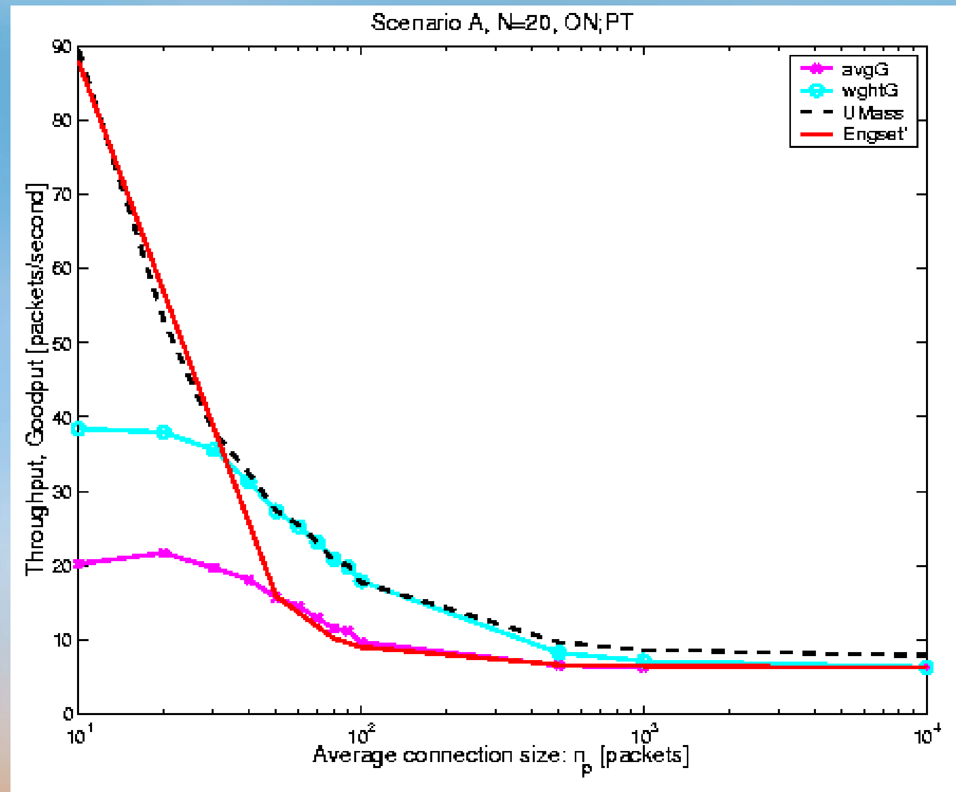

- **All models over-estimate throughput for short connections (initial slow-start)**
- **Good fit for n_p>40**

Comparison (IV): Scenario (B)

Variation of # Sources, N

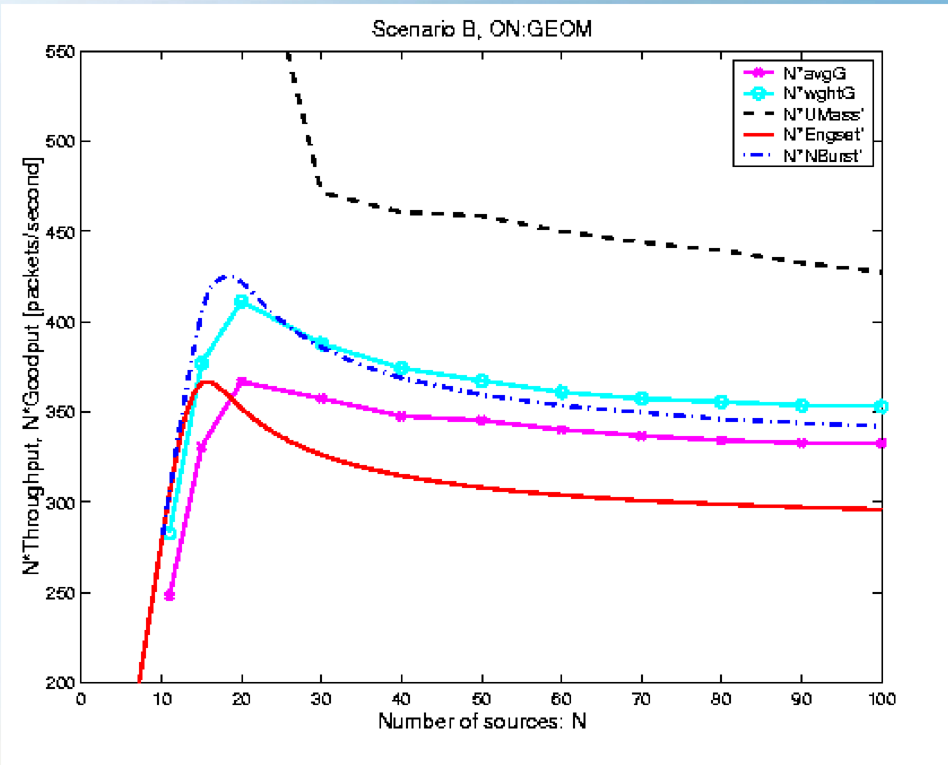

Variation of Buffer-size, B

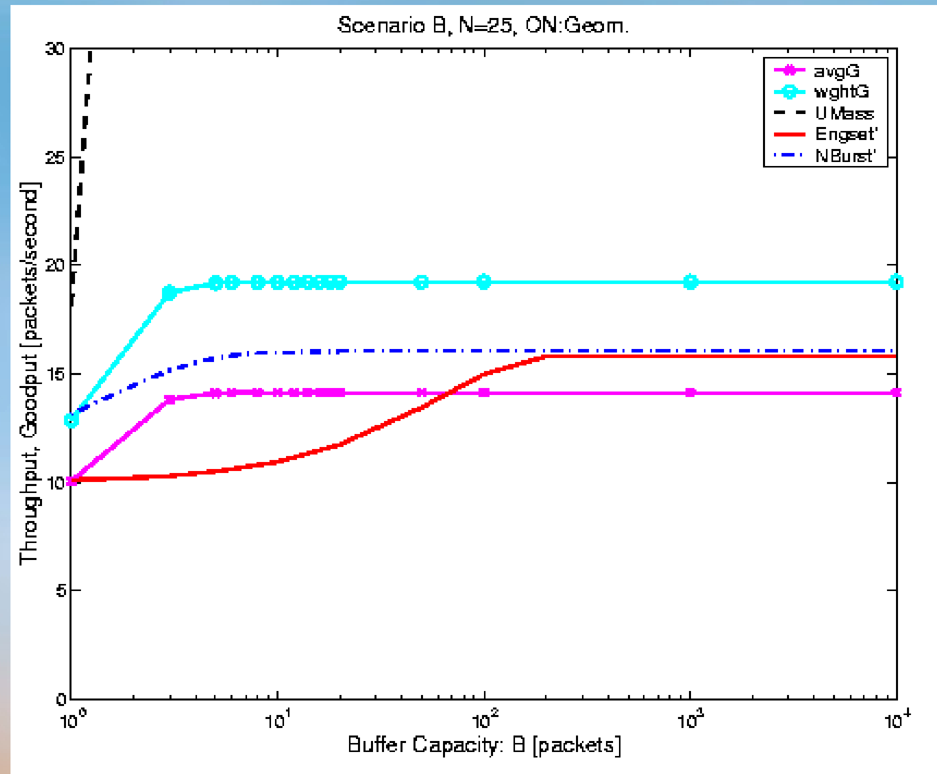

- **Best fit of TCP-NBurst**
- **TCP-UMass over-estimates largely**

- **Goodput insensitive too large buffers**
- **TCP-UMass model useless**

Summary & Outlook (Wireless scenarios)

Comparison of analytic models, TCP-UMass, TCP-Engset, and TCP-NBurst, with simulation of TCP Reno:

- TCP-UMass model requires additional (performance) parameters as input
- All models fail for short connections (initial slow-start not regarded)
- PT distributed connection sizes cause gap:
  weighted average >> plain average of goodput per connection

... More experiments are required

Outlook:

+ Include properties of wireless link
  - Long delays (d_3)
  - Error model / packet-loss model for wireless link

  → Necessary: Extension of analytic models for wireless links

+ Investigate impact of TCP enhancements: SACK, SNOOP