TCP Behavior in Networks with Dynamic Propagation Delay

Mark Allman, NASA GRC/BBN

Jim Griner, NASA GRC

Alan Richard, NASA GRC/Analex

IEEE Globecom

November, 2000

Overview

- Introduction and Motivation
- Simulation Environment
- RTO Estimator Validation
- Simple TCP Experiments
- Simple Handoff Experiments
- Conclusions

Introduction and Motivation

- Plenty of researchers have looked at the impact of long, static delays on TCP performance.
  - See RFCs 2488, 2760 and references therein.
- But, what about situations where the propagation delay changes over time?
  - E.g., NASA's Earth-observing satellites.

Introduction and Motivation (cont.)

- Our paper is based on models of satellites sending data to the ground.
- However, we believe the results apply to any situation where modest motion is involved.

Simulation Environment

- We used a variety of spacecraft orbiting in the LEO and MEO bands.
  - These spacecraft send data to TDRS, which transmits the data to Earth.
- We used Satellite Toolkit 4.0 to generate orbital data.
- We introduced a variable delay link into the *ns* network simulator.
  - The propagation delay along the link changes as a function of time, based on the STK output.

Simulation Environment (cont.)

- Simulated topology:

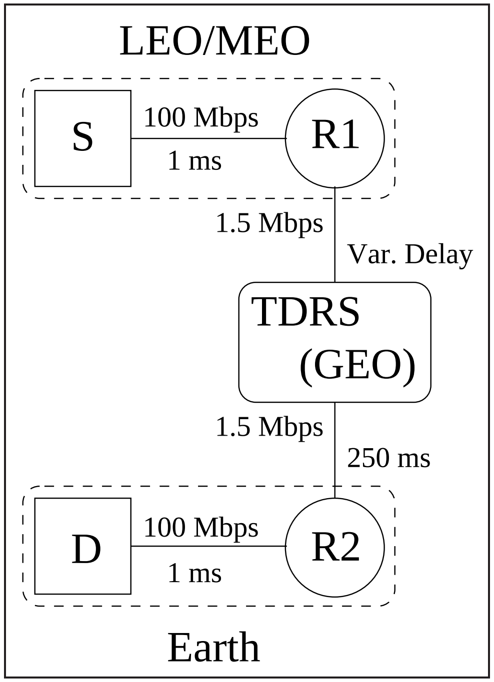

Variable Delay Scenarios

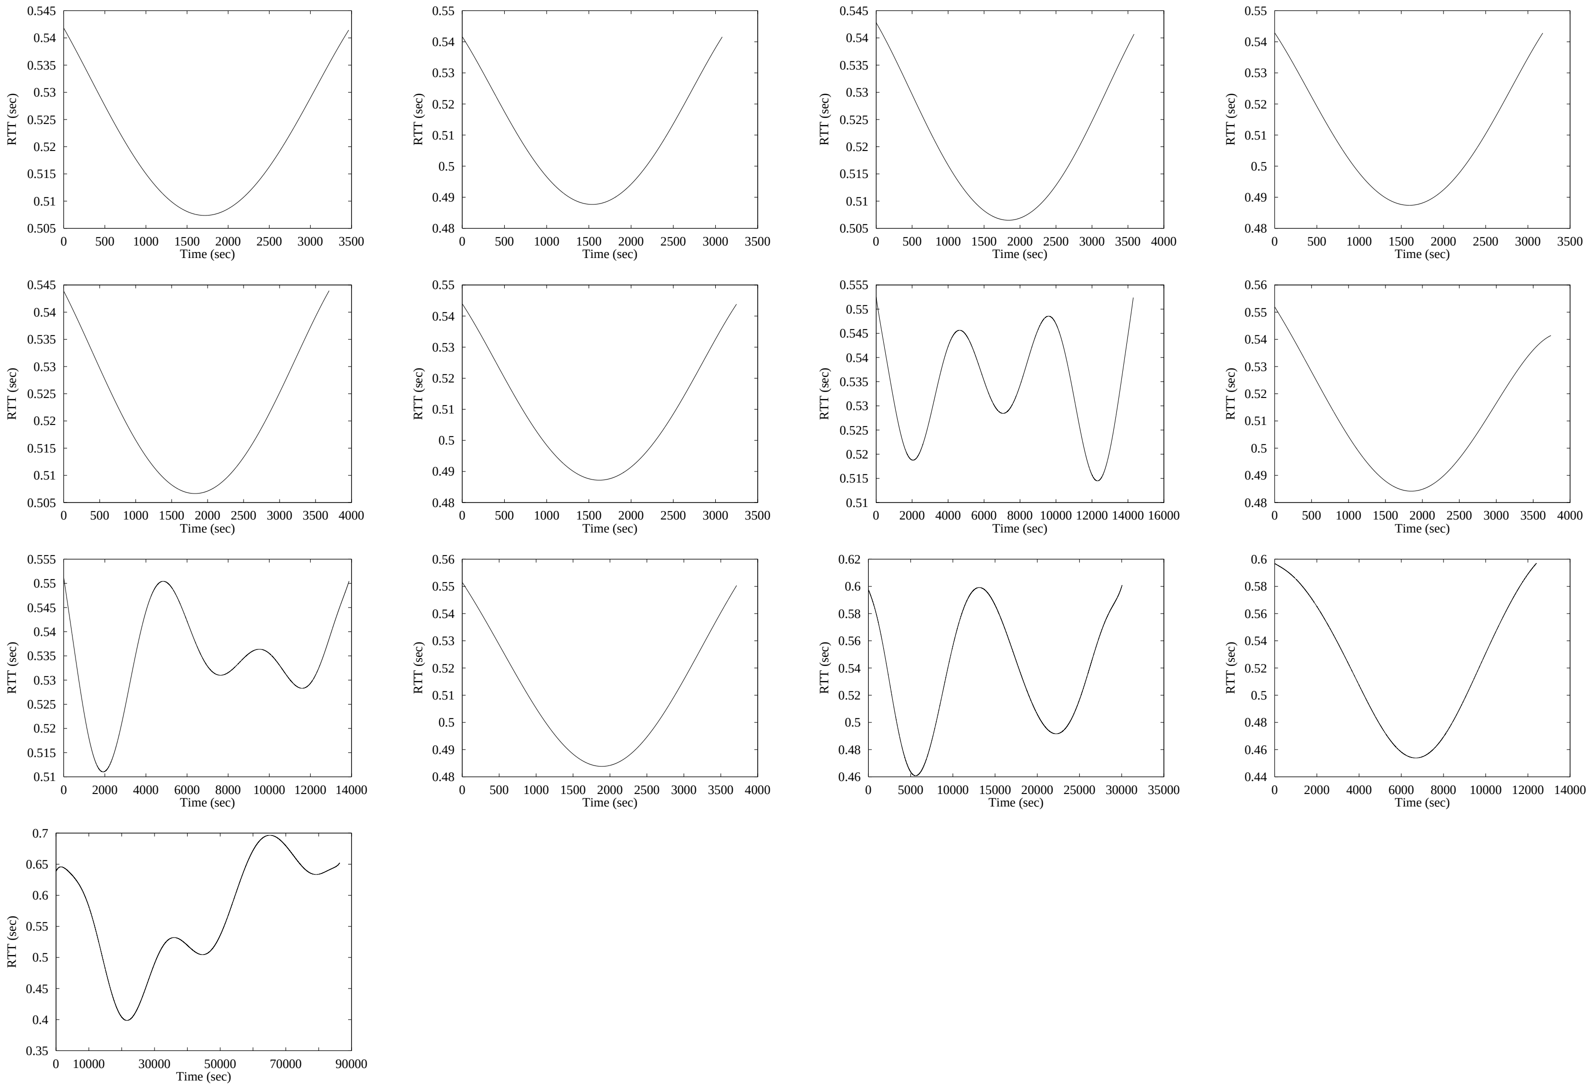

Simple RTO Experiments

- TCP uses a *retransmission timer* (RTO) to guarantee reliable data delivery.
- The standard RTO estimator:

$$RTO \leftarrow SRTT + 4 \cdot RTTVAR$$

- RTO measured and calculated using a clock with granularity G.
  - Traditionally $G = 500\ ms$
  - Some have suggested finer grained timers will yield better performance, so we also used $G = 1\ ms$.

Simple RTO Experiments (cont.)

- Loss is also taken as an indication that the network is congested.
  - Hence, the sending rate is reduced.
- Therefore, one desirable property of an RTO estimator is that it not retransmit segments too early and cause a needless reduction in sending rate.

Simple RTO Experiments (cont.)

- Do the variable delay scenarios used in our experiments confuse the RTO estimator?
  - Set the maximum TCP window size to 1 segment.
  - Run a TCP transfer for the length of the scenario.
  - Watch for retransmissions.
- Answer: No. The RTO estimator is able to cope with the changing propagation delays we tested.
  - But, what about a slightly more dynamic environment with queueing delays?

Single Flow Tests

- Tested various file sizes (4–10,000 packets).
- The transfer start time was roughly every 60 seconds over the course of the scenario.
- Started with $G = 500\ ms$

Single Flow Tests (cont.)

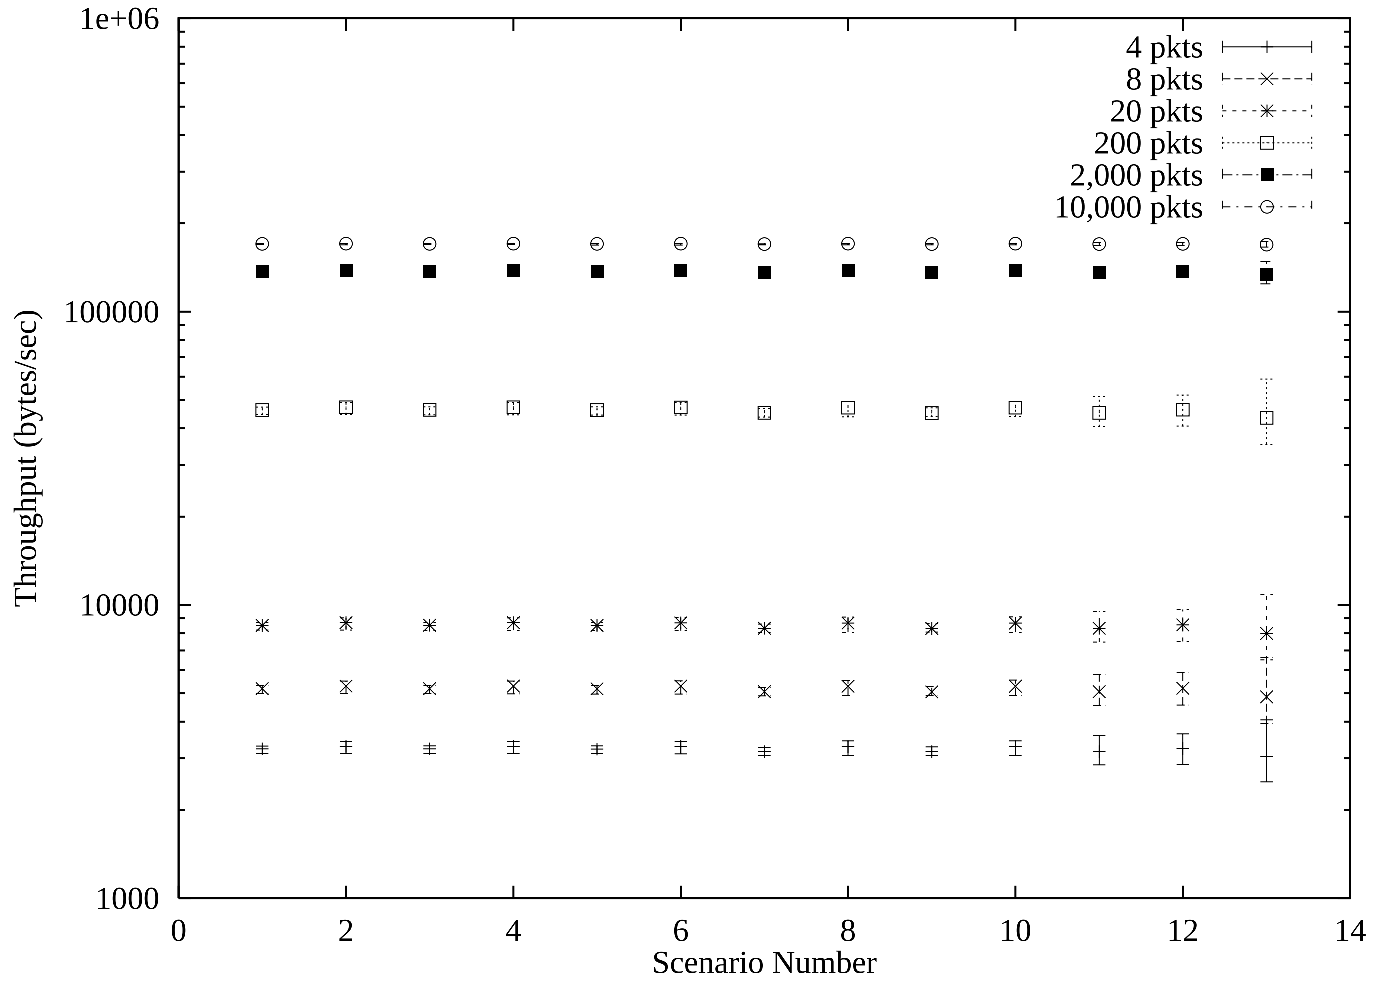

Single Flow Tests (cont.)

- As expected...
  - Small files underutilize the capacity.
  - Large files nearly fully utilize the capacity.
  - More throughput variation in small files.
- Also, no unnecessary retransmits were detected.

Single Flow Tests (cont.)

- What about using a fine-grained timer?
  - Small transfers (4–200 packets) did not cause needless retransmissions.
    - Small transfers do not build queues – and we know that fine-grained timers work well with no queues on our delay scenarios.
    - $RTTVAR$ is initially $\frac{RTT_{meas}}{2}$, which inflates the RTO at the beginning of a transfer, providing some protection against spurious retransmits.
  - Large transfer *do* experience needless retransmits.

Single Flow Tests (cont.)

- 2,000 packet transfer

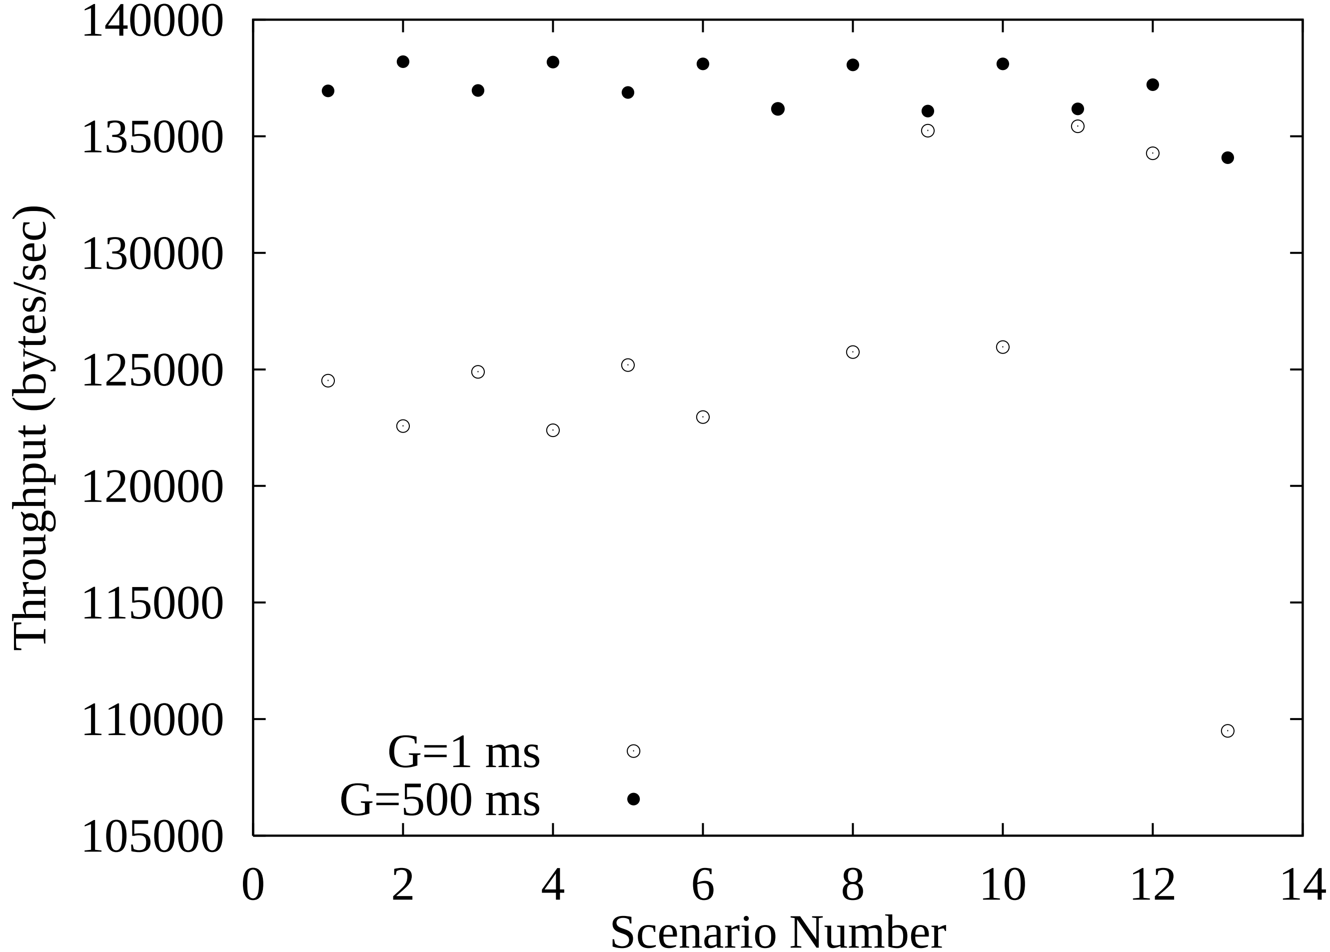

Handoff Scenario

- Our last scenario models a perfect (no loss, no reordering) handoff that essentially moves from a single GEO hop to a double hop and back.
- $G = 1\ ms$ cannot cope with the drastic change in RTT caused by moving from a single hop to a double hop.
- $G = 500\ ms$ does not needlessly retransmit even when crossing the large jump in throughput.

Conclusions

- With a large minimum RTO (e.g., as we get with $G = 500\ ms$) TCP performs quite well in the environments examined.
- Fine-grained timers reduce performance for long transfers.
- As in more static environments, short transfers often underutilize the capacity of the network path.
- The throughput obtained by short transfers is somewhat variable depending on start time.

Future Work

- Consider more realistic handoffs where reordering and/or loss may occur.

- When a satellite is moving, typically the signal strength is changing, as well as the propagation delay. This will yield different BERs at different points in the curve. This should be investigated.

- A more realistic traffic pattern should be obtained and used.