Quick-Start for TCP and IP

Draft-amit-quick-start-03.txt

A. Jain, S. Floyd, M. Allman, and

P. Sarolahti

ICIR, December 2004

www.icir.org/floyd/talks/quickstart-Dec04.pdf

www.icir.org/floyd/talks/quickstart-Dec04.ppt

QuickStart with TCP, for setting the initial window:

- In an IP option in the SYN packet,
  the sender's desired sending rate:
  - Routers on the path decrement a TTL counter,
  - and decrease the allowed sending rate, if necessary.
- The receiver sends feedback to the sender in the SYN/ACK packet:
  - The sender knows if all routers on the path participated.
  - The sender has an RTT measurement.
  - The sender can set the initial congestion window.
  - The TCP sender continues using normal congestion control..
- From an initial proposal by Amit Jain

The Quick-Start Request Option for IPv4

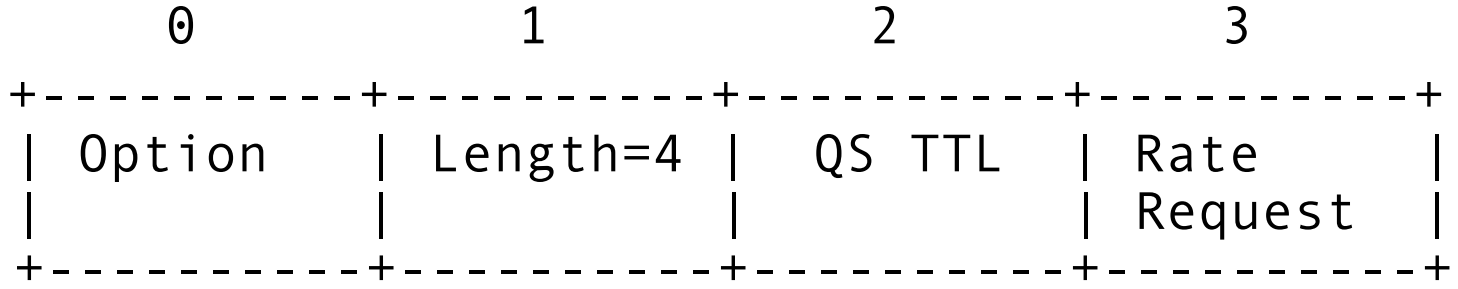

- Explicit feedback from all of the routers along the path would be required.
- This option will only be approved by routers that are significantly underutilized.
- No per-flow state is kept at the router.

Quick-Start in the NS Simulator

- Added to NS by Srikanth Sundarrajan.

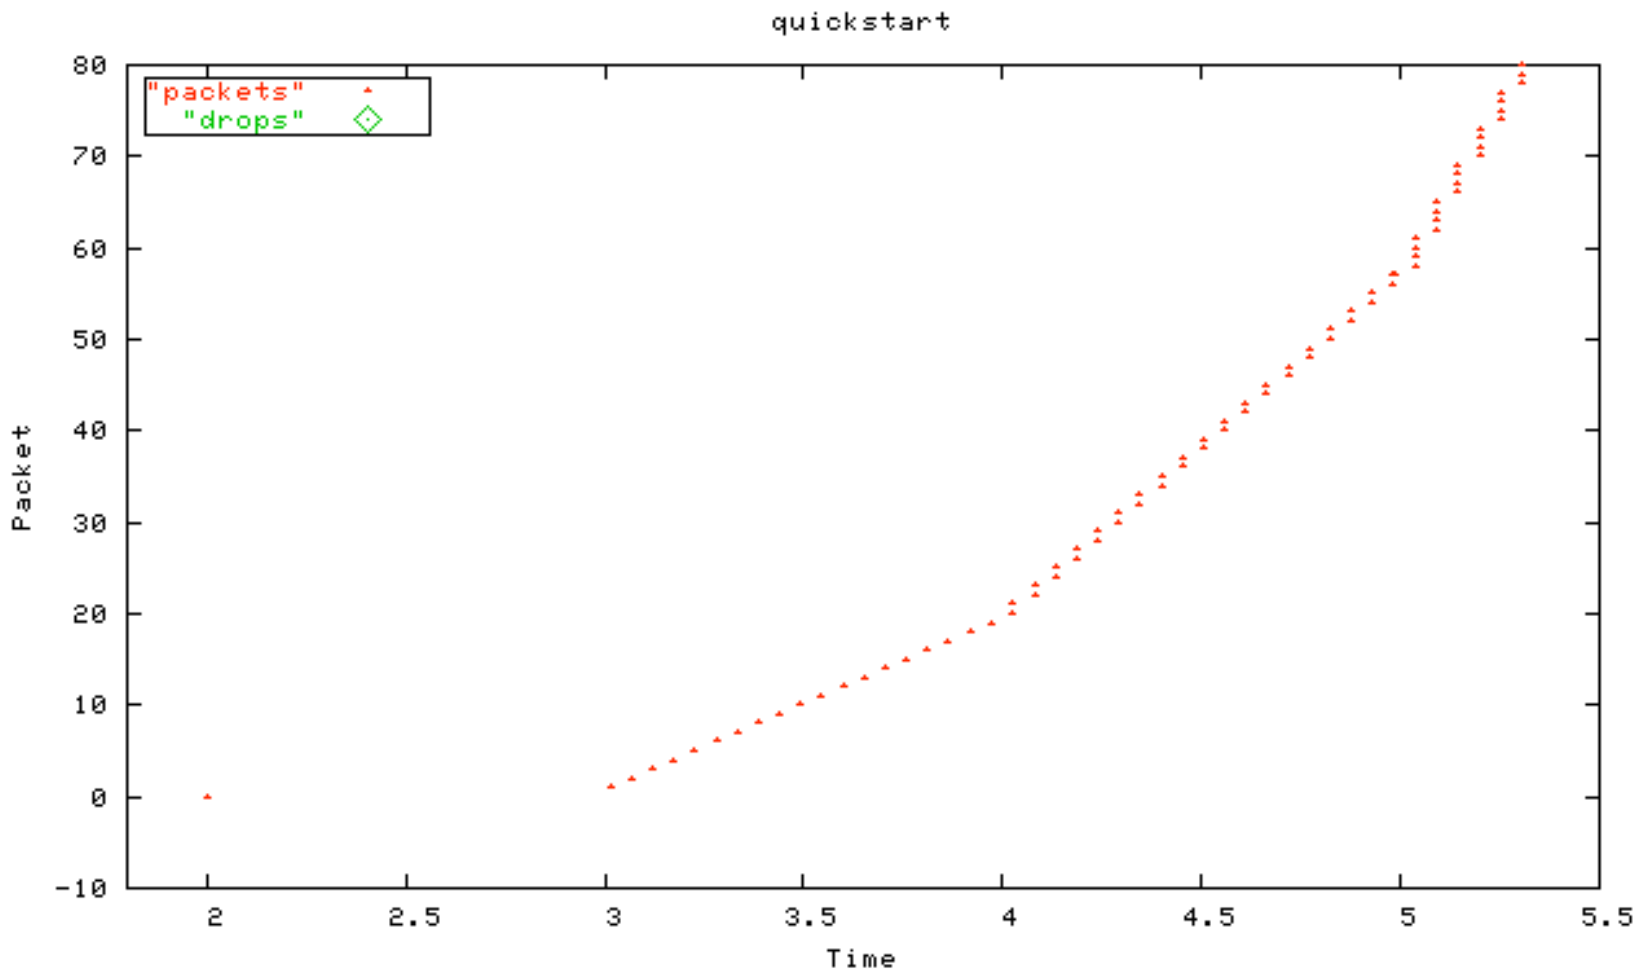

A Failed QuickStart Request:

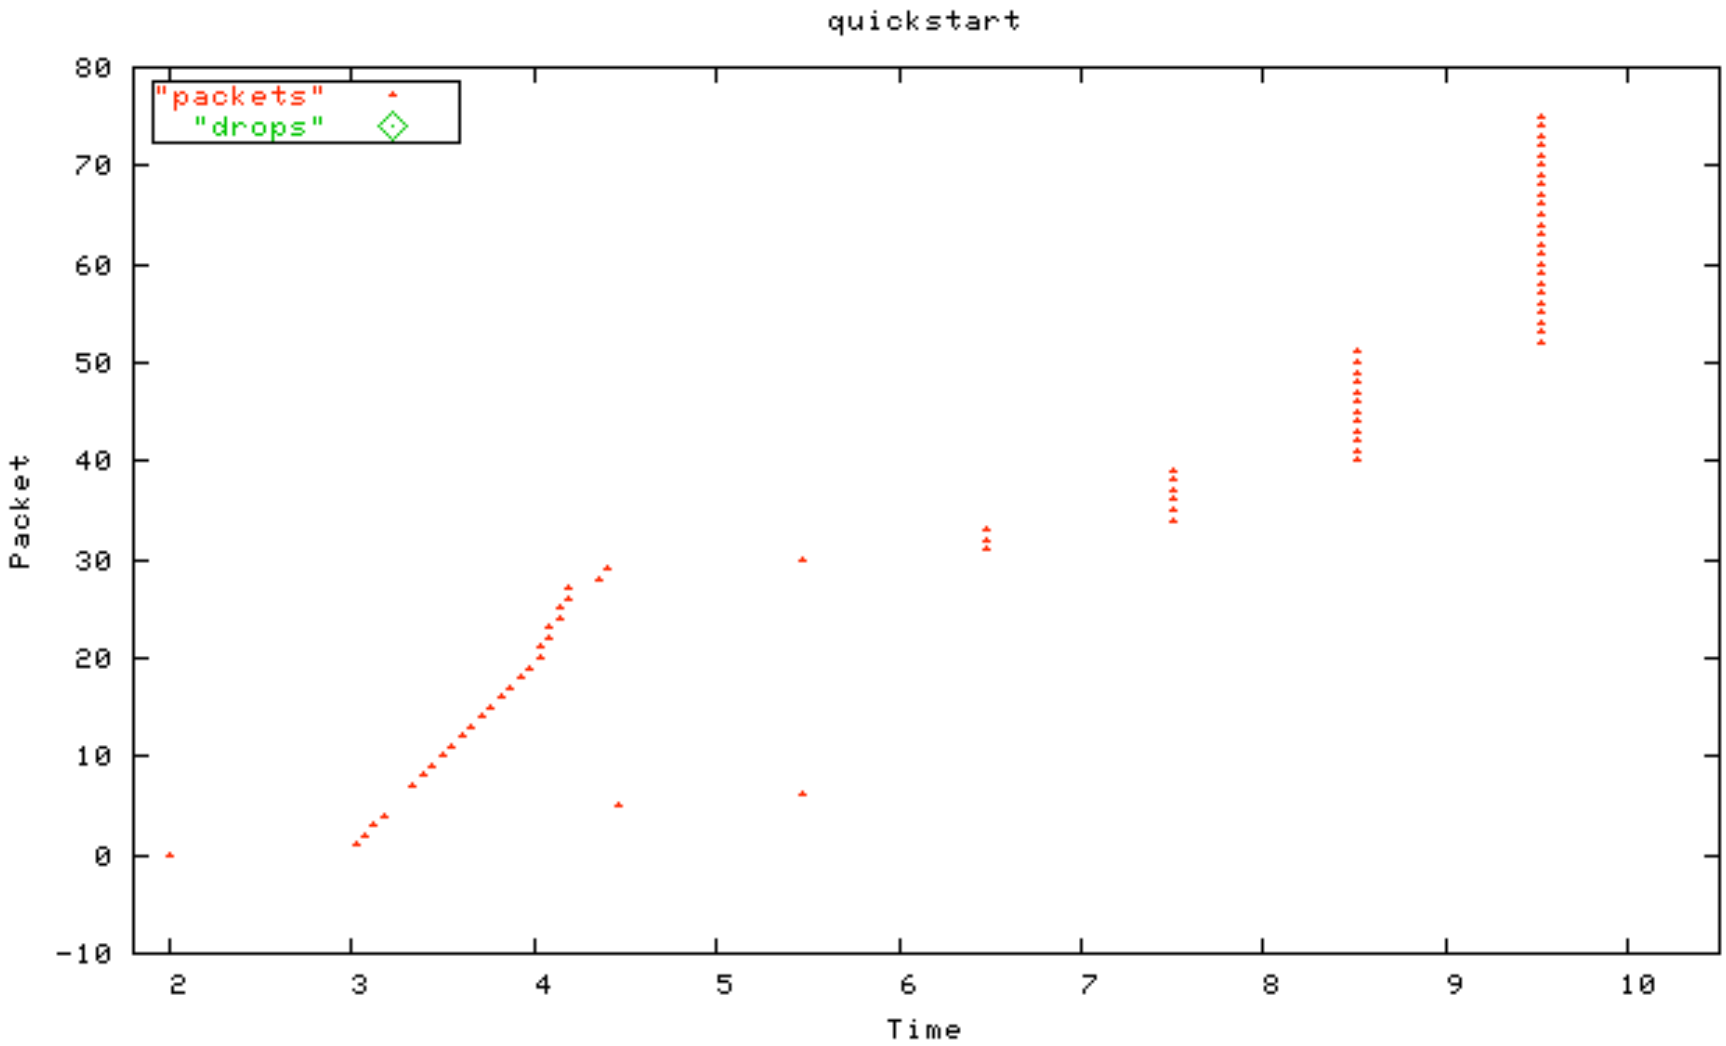

Changes from draft-amit-quick-start-02.txt:

- Using Quick-Start in the Middle of a Connection.
  - The request would be on the total rate, not on the additional rate.
  - Examples: after idle periods, mobility events.
- The request is now in bytes per second, not packets per second.
- New sections include:
  - When to use Quick-Start (QS)
  - TCP: Responding to a Loss of a QS Packet
  - Quick-Start with DCCP
  - Quick-Start in IP Tunnels

What would be needed in routers?

- T: Configured QuickStart threshold (in Bps).
  - Requires knowledge of output link bandwidth.
- L: Current link utilization (in Bps).
- R: Recent granted QuickStart requests (in Bps).
  - Requires state of aggregate of granted requests.
- Router algorithm:
  - Max request to grant: T - L - R Bps

Possible Initial Deployment Scenarios:

- Intranets:
  - Centralized control over end nodes and routers.
  - Could include high-bandwidth, high-delay paths to remote sites.
- Paths over satellite links:
  - High bandwidth, high delay
- 2G/3G networks:
  - RTTs of up to one second

Design Issues: IP Options, ICMP, or RSVP?

- IP Options:
  - Blocked by some middleboxes
  - Takes the slow path in routers?
- ICMP:
  - Blocked by some middleboxes.
  - Mechanisms would be needed to address all routers along the path, and get one answer.
- RSVP:
  - Soft state in routers is not needed.

Design issues: the encoding of the Rate Request

- Linear function:
  - Minimum request of 80 Kbps, maximum request of 255*minimum, or 20.5 Mbps.
  - 80-Kbps increments
- Powers of two:
  - Use 4 bits of the 8-bit field.
  - Minimum request of 80 Kbps, maximum request of 2^14*minimum, or 1.3 Gbps.
  - Doubling the request from one to the next.

Related Work:

- Faster Start-up without modifying routers:
  - Packet-pair and extensions.
- Less-than-best-effort for the initial window.
- Other forms of feedback from routers:
  - Free buffer size, available bandwidth.
- New congestion control mechanisms.
  - E.g., XCP, AntiECN.

Questions:

- Would Quick-Start be deployable?
  - Given the chicken-and-egg problems (requiring deployment in both routers and end nodes).

Would QuickStart be deployable, given the tighter integration required across the Internet?

- End nodes can make spurious requests.
- Routers can grant requests inappropriately.
- Colluding access routers can cheat (though the benefit would be minimal).
- It would be harder to police end-nodes.
- Problems from middleboxes.
- Problems with congestion at non-IP queues.
- Additional delay for the QuickStart Request
  - (if it doesn't take the fast path in routers).

Questions:

- Is something like this needed?
  - Are there going to be underutilized paths?
- Would the benefits of Quick-Start be worth the added complexity?
  - Quick-Start Request packets would not take the fast path in routers.
- What would be the relationship between Quick-Start and new router-based congestion control mechanisms (e.g., XCP)?