TFRC for Voice: the VoIP Variant

Sally Floyd, Eddie Kohler.

March 2005

http://www.icir.org/floyd/papers/

draft-ietf-dccp-tfrc-voip-01.txt

http://www.icir.org/floyd/talks.html

First, four repeat viewgraphs from last IETF…

VoIP: fairness in Bps.

- Standard TFRC has the goal of fairness in pps with TCP flows using the same packet size.

- The VoIP variant of TFRC has the goal of fairness in Bps with TCP flows using 1500-byte packets, (following RFC 3714).

- The VoIP variant assumes optimistically that the network limitation is in Bps, not in pps.

VoIP: fairness in Bps.

- In the TCP throughput equation, use the measured loss event rate and a packet size of 1460 bytes.

- Reduce the allowed transmit rate to account for the fraction of the VoIP bandwidth that would be used by 40-byte headers:
  - X <- X * TruePktSize/(TruePktSize + Header)
    - TruePktSize = average segment size in bytes
    - Header = 40 bytes

- Enforce a Min Interval between packets of 10 ms.

Measuring Congestion:

- The VoIP variant of TFRC uses the loss event rate.
  - RFC 3714 uses the packet drop rate.
- These are both affected by packet size and by the smoothness of the sending rate.
- The effect of packet size on the packet drop rate could use more investigation.

The VoIP variant of TFRC:

- As it stands now, it sometimes favors the VoIP TFRC flow over the large-packet TCP flow.

- This needs to be quantified and evaluated.

Changes from draft-ietf-dccp-tfrc-voip-00.txt:

- Added more simulations.
- Added a Related Work section.

Small Packets, Standard TFRC:

- The problem:
  - TCP flows with 1460-byte packets, competing against:
  - Standard TFRC flows with 200-byte packets, 10 ms between packets (e.g., 160 Kbps)

Small Packets, Standard TFRC:

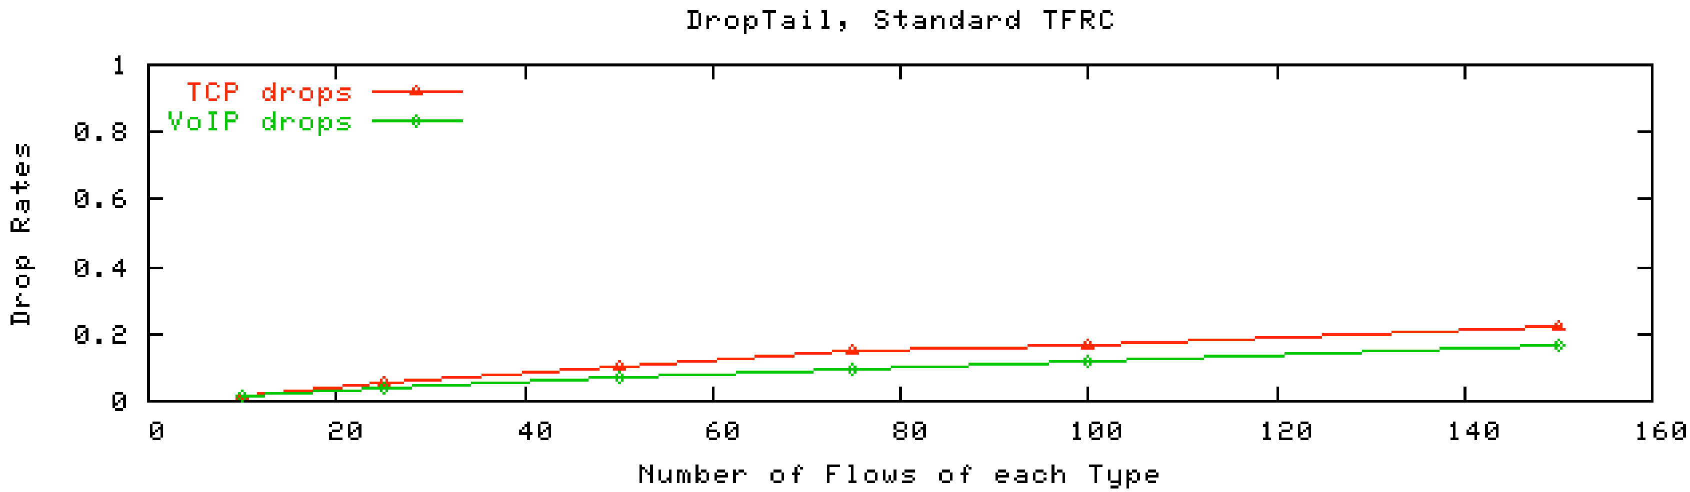

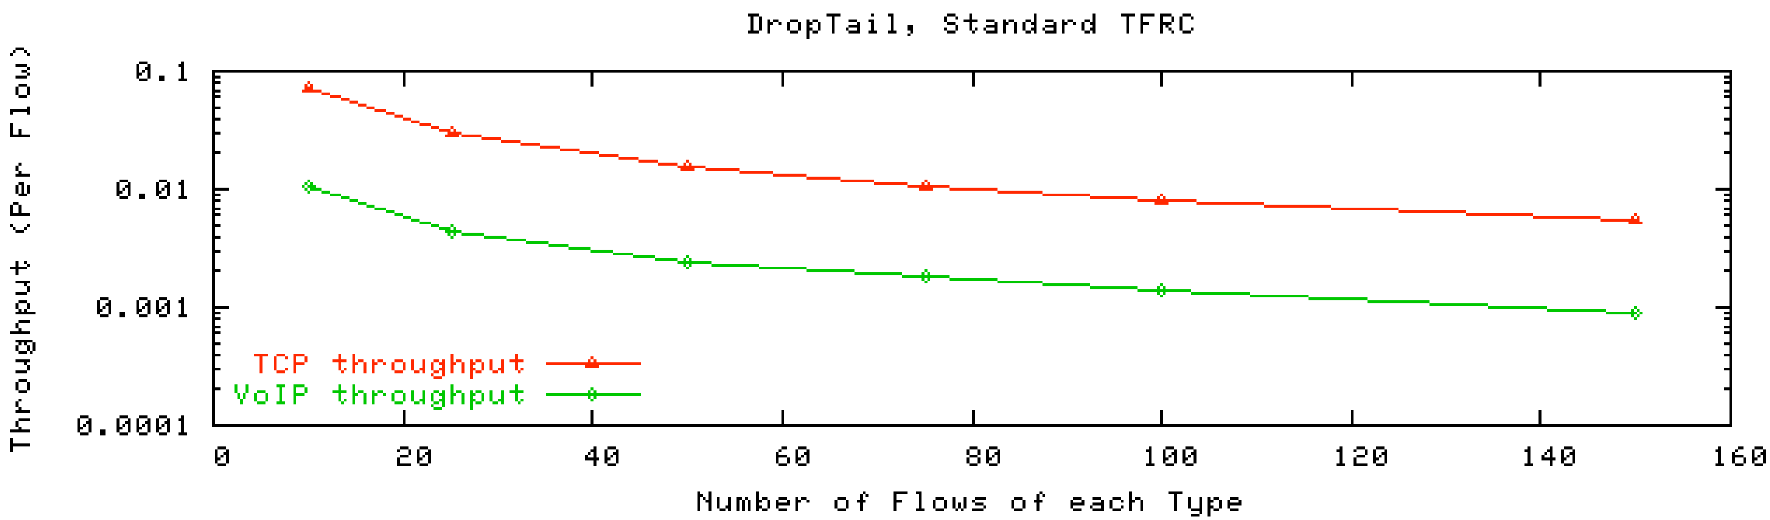

Simulations with VoIP TFRC:

- The algorithm so far in the draft:
  - TCP flows with 1460 bytes, competing against
  - VoIP TFRC flows with 200-byte packets.

Simulations with VoIP TFRC:

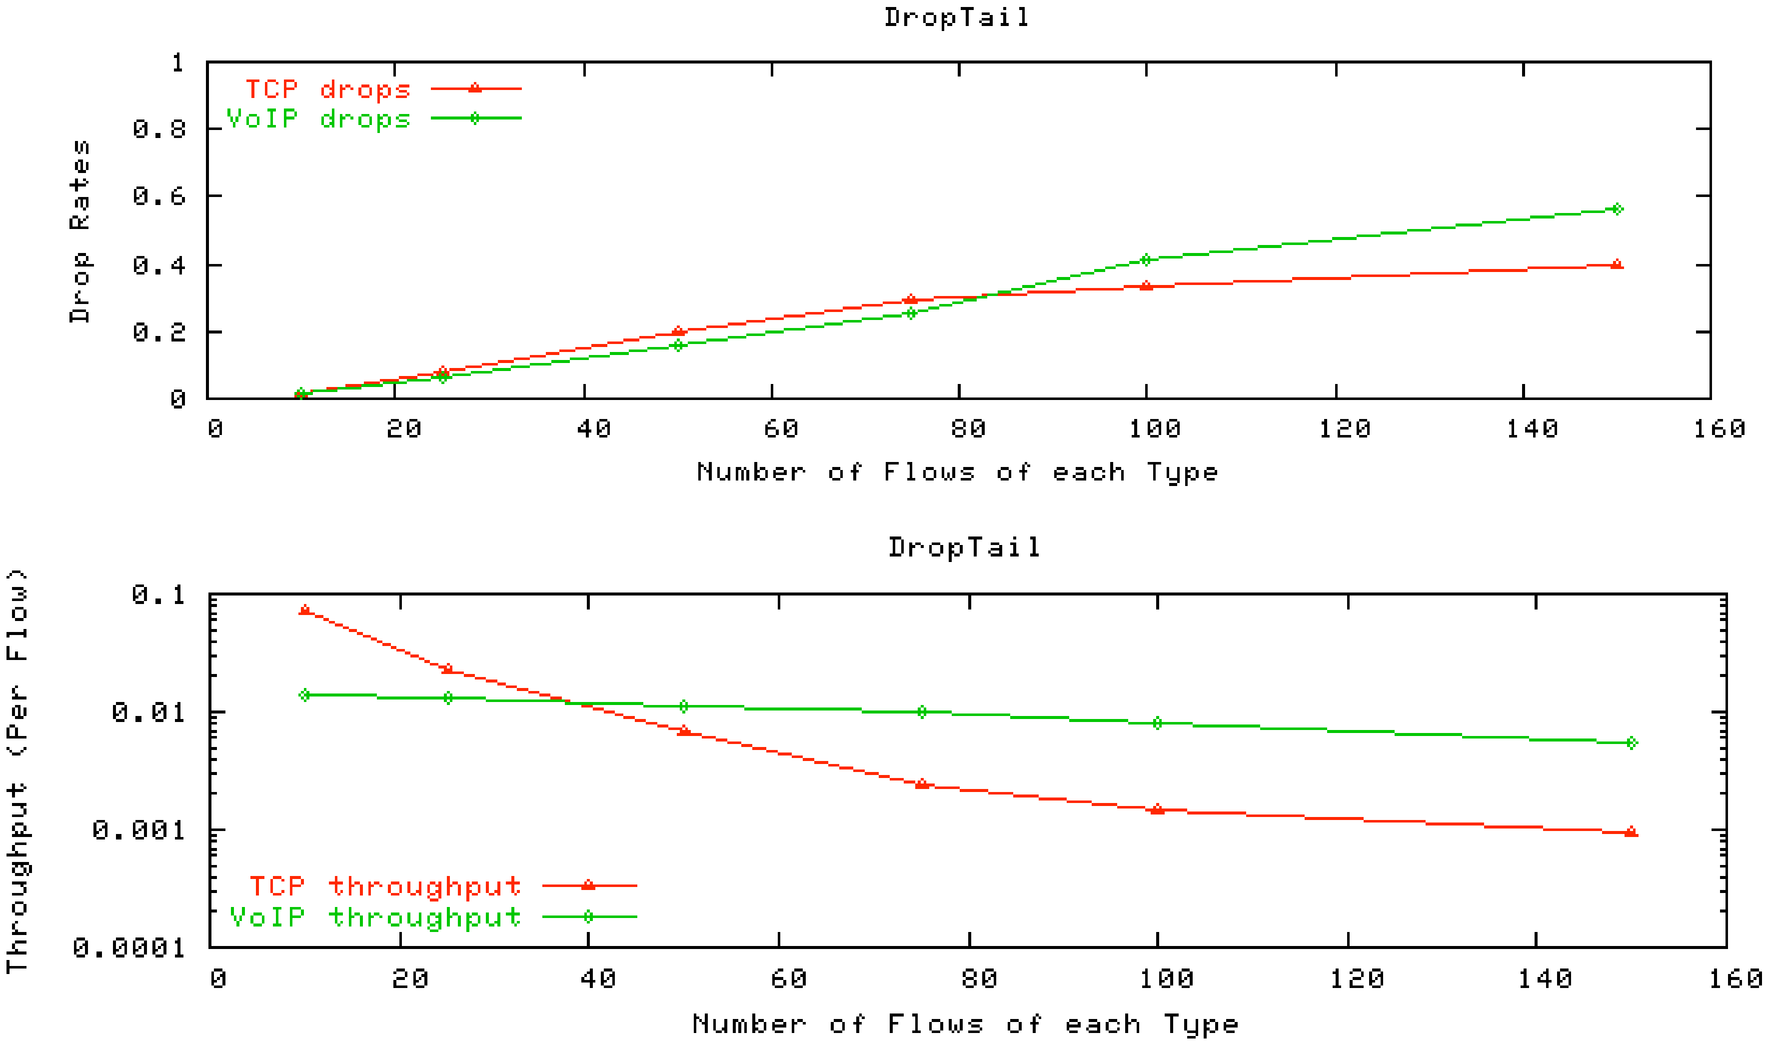

Table 1 from draft-ietf-dccp-tfrc-voip-01.txt.

This won’t do.

- The problem [Widmer, Boutremans, and Le Boudec, 2004; Vasallo 2000]:
  - TCP and TFRC can’t send multiple packets per RTT in the face of high congestion; VoIP TFRC can.
  - The loss event rate used by TFRC follows TCP by responding to at most one loss per RTT.
  - When TCP has one “loss event” for every two 1460-byte packets, a rate-equivalent TFRC VoIP might have one “loss event” for every 30 100-byte packets.
  - So the “loss event rates” computed by VoIP TFRC can be too low in times of high congestion.
- One solution:
  - Loss rates have to be calculated differently, at least in times of high loss.

Solutions:

- Previously-explored solutions [WBB04]:
  - Count virtual packets [V00].
  - Do random sampling of arriving packets.
  - Shorten part of the Loss Interval.
- One possible solution for VoIP TFRC:
  - For short loss intervals (at most two RTTs), count the actual packet loss rate (but don't increase the number of loss intervals).
  - Somewhat like "random sampling" above, but only for short loss intervals.
  - Examples: "./test-all-friendly HighLossShort" and "./test-all-friendly printLossesShort" in tcl/test in NS.

The Modified TFRC VoIP:

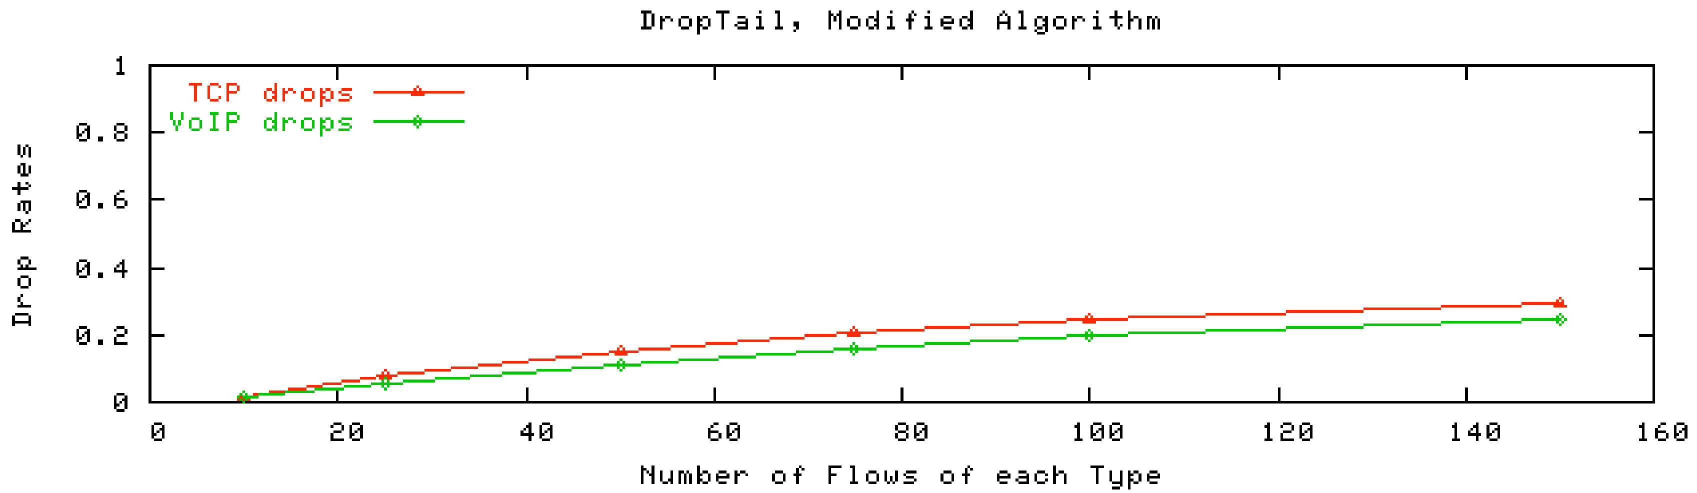

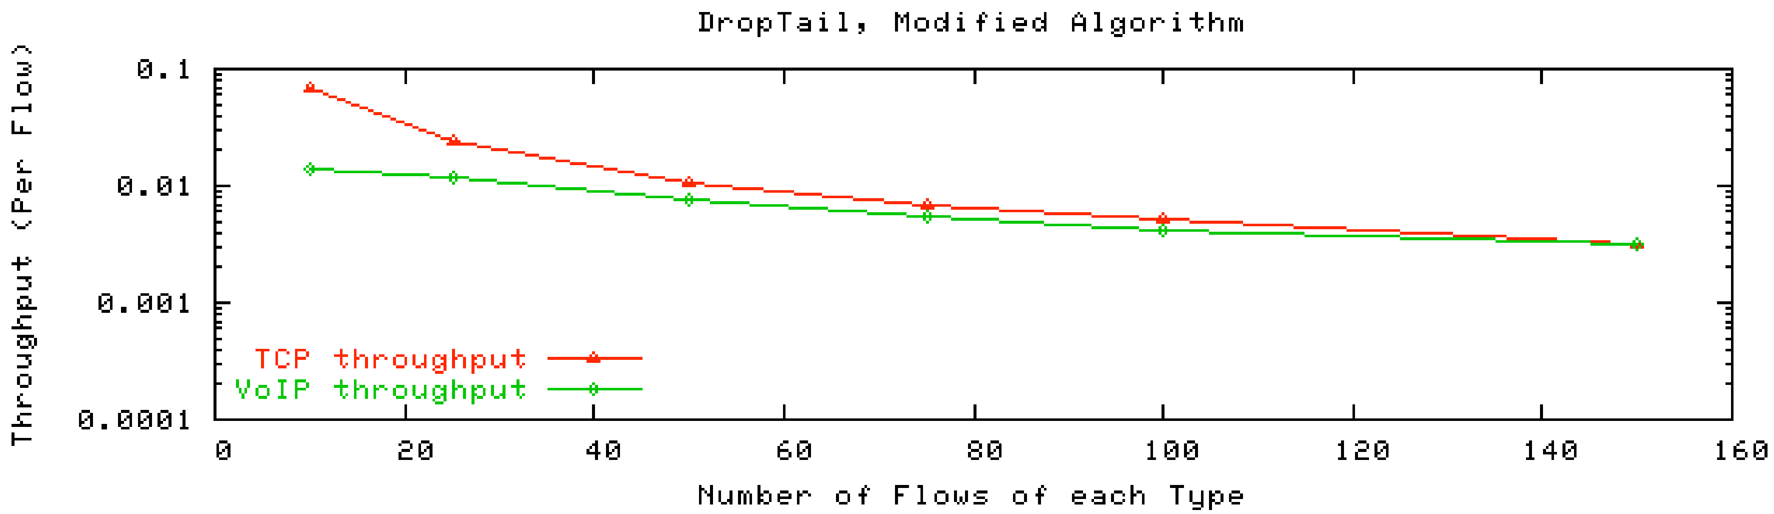

Issues remaining:

- More exploration needs to be done.
- There were two problems:
  - VoIP TFRC can send many packets in one RTT, even in the face of heavy congestion. Taken care of.
  - VoIP TFRC, with small packets, sees different packet drops that it would have with larger packets. When is this a problem?
- For simulations with RED in byte mode (where small packets are less likely to be dropped than large packets):
  - Even the modified VoIP TFRC gets much more than its share of the bandwidth in times of high congestion.
  - Under investigation.

Assumptions:

- The sender doesn’t know the packet-marking mechanisms used by the routers.
  - E.g., Drop-Tail? Queue in bytes or in packets? AQM in byte mode or in packet mode?
- The sender can’t assume that packets or bytes are being dropped with some relatively stable dropping probability p.
  - This is not necessarily the case.

Assumptions:

- The number of packets dropped in one round-trip time is not necessarily an indication of the level of congestion.

Extra Viewgraphs:

Measuring Congestion:

- Packet size in a Drop-Tail world:
  - Queue measured in bytes, packets, or in-between?
  - Smooth or bursty sending rates?
  - High or low levels of statistical multiplexing?
- RED in packet mode:
  - Same packet drop rate for big and small packets.
  - TFRC measures the loss interval in packets.
- RED in byte mode:
  - Same byte drop rate for big and small packets.

The state of TFRC in NS:

- Includes the VoIP variant.
- Includes RFC 3390 initial sending rates.
- More updating is needed.
  - Add RFC 3390 sending rates after idle periods.
  - Add Faster Restart.
  - Add overhead for packet headers.