Controlling High Bandwidth Aggregates

Sally Floyd
Joint work with Steve Bellovin, Ratul Mahajan, Vern Paxson, Scott Shenker
November 29-30, 2000, E2E Research Group

Overview of Talk:

- Controlling High-Bandwidth Flows at the Congested Router
- Local Aggregate-based Congestion Control
- Pushback

Controlling High-Bandwidth Flows at the Congested Router

Ratul Mahajan and Sally Floyd, http://www.aciri.org/floyd/papers/

- RED-PD (RED with Preferential Dropping)

- Uses the RED packet drop history to detect high-bandwidth flows.

- Packets from monitored flows are preferentially dropped before the output queue.

- Flows: defined by IP source/destinate address and port numbers (or by Security Association, for IPsec).

- Monitored flows: either nonconformant, or conformant flows with small round-trip times.

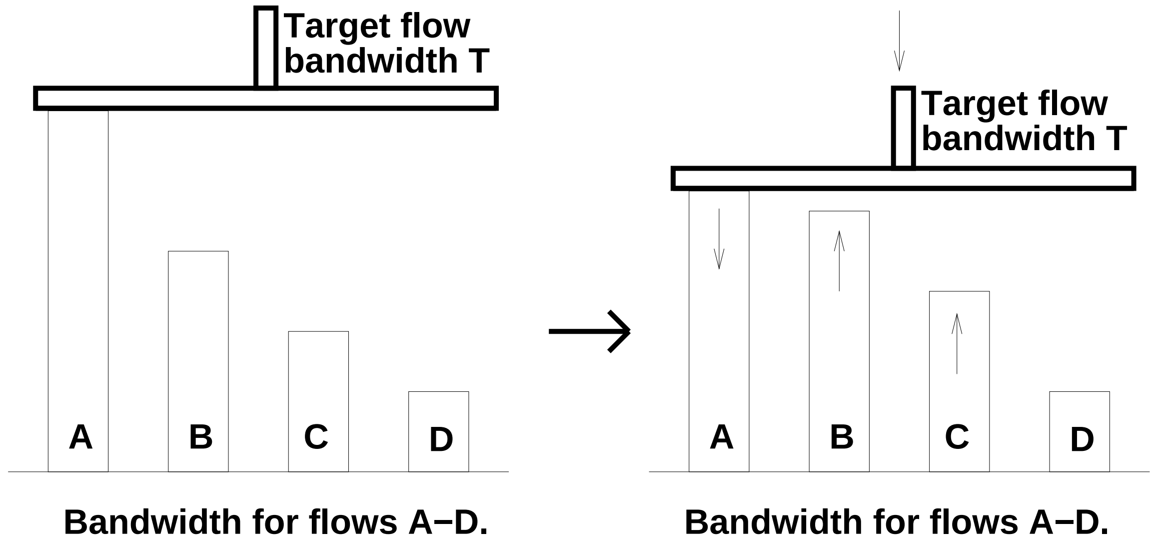

Restricting flows to a target bandwidth T.

- The target bandwidth T is $\frac{\sqrt{1.5}}{R\sqrt{p}}$.
  R: a configured round-trip time
  p: the current packet drop rate

.

- After flows are preferentially dropped, identifying non-conformant flows would be a separate step.

Controlling High-Bandwidth Aggregates

- Similarities between controlling aggregates and flows:
  - Both use the packet drop history for identification.
  - Both use preferential dropping before the output queue.

- Differences:
  - Aggregate-based congestion control (ACC) should rarely be invoked.
  - Aggregates can have fuzzy, overlapping definitions.
  - There is no simple fairness goal for aggregates, as for flows.

A Thought Experiment of Aggregate-based Congestion Control (ACC):

- No flash crowds:
  - N aggregates A_1-A_n share link with background traffic.
  - Packet drop rate p (e.g., $p = 0.01$).

- Flash crowd i from aggregate A_i, no ACC:
  - During the flash crowd, the drop rate is p_1 (e.g., $p_1 = 0.2$).
  - The throughput for A_j, for $j \neq i$, is roughly $\frac{1}{\sqrt{p_1/p}}$ of its value without the flash crowd (e.g., 1/5-th of its old value).

- Flash crowd i with ACC:
  - If during A_i's flash crowd, A_i is restricted to at most half the link bandwidth:
  - A_i's throughput is at worst halved, compared to no ACC.
  - All other traffic has its throughput at worst halved, compared to no flash crowd, and therefore its packet drop rate at most quadrupled.

The Mechanisms of Aggregate-based Congestion Control:

- Detect sustained congestion, as characterized by a persistent, high packet drop rate.

- Look at the packet drop history:
  – See if the packet drops are heavily represented by some aggregate (e.g., as defined by destination address prefix, source address prefix, etc.).

- If an aggregate is found:
  – Preferentially drop packets from the aggregate before they are put in the output queue, to rate-limit aggregate to some specified bandwidth limit.

Now consider a Distributed Denial of Service (DDOS) Attack:

- If an aggregate causing congestion is from a DDOS, then the aggregate will contain both malicious traffic and legitimate, ”good” traffic.

- Because of spoofing, we can not necessarily trust the IP source addresses.

- ”Pushing-back” some of the preferential dropping of the aggregate to neighboring, upstream routers:
  - Does not rely on valid IP source addresses.
  - Limits the damage from the DoS attack, reducing wasted bandwidth upstream.
  - In some cases, allows preferential dropping to be concentrated more on the malicious traffic, and less on the good traffic within the aggregate.

Pushback, Traceback, and Source Filtering:

- With Pushback, a router rate-limiting packets from aggregate A might ask upstream routers to rate-limit that aggregate on the upstream link.

- Pushback is orthogonal to ”traceback”, which tries to trace back an attack to the source.
  - Traceback allows legal steps to be taken against the attacker.
  - Traceback is of limited effectiveness in a highly distributed attack.

- Pushback is orthogonal to source filtering, which limits the ability to spoof IP source addresses.
  - Source filtering is important in any case.
  - DoS attacks can also come from valid source addresses.
  - Pushback can be useful even when source addresses can be trusted.

Questions about Aggregate-Based Congestion Control:

• How often do routers have periods of sustained, high packet drop rates?

• For periods of high packet drop rates, how often is it due to:
(1) DOS attacks? (Local ACC and pushback would help.)
(2) Legitimate flash crowds? (Local ACC would help, pushback would be OK.)
(3) Network problems (e.g., routing failures)?
(4) Diffuse general congestion?
For (3) and (4), ACC will be of limited effectiveness, and probably won't be invoked.

• Would ACC for legitimate flash crowds be a useful incentive for web servers to use effective web caching and/or content distribution?