Aggregate-Based Congestion Control and Pushback

Sally Floyd
Joint work with Steve Bellovin, John Ioannidis, Ratul Mahajan,
Vern Paxson, Scott Shenker

Dec. 5, 2000, ACIRI Annual Review

Overview of talk:

- What are the problems?
  – Bullies, mobs, and crooks.
- Controlling misbehaving or high-bandwidth flows (i.e., bullies).
- Controlling flash crowds (i.e., mobs).
- Controlling Denial-of-Service attacks (i.e., crooks).

What are the problems?

Bullies (misbehaving or high-bandwidth flows):

- Flow: defined by source/destination IP addresses and port numbers.
  - Example: a single TCP connection.

- Problem: Fairness between competing flows.

- Problem: Preventing congestion collapse.
  - From congested links carrying undelivered packets .
  - Floyd, S., and Fall, K.,

"Promoting the Use of End-to-End Congestion Control in the Internet", IEEE/ACM Transactions on Networking, August 1999.

What are the problems?
Mobs (flash crowds):

- Example: The Starr Report, September 11, 1998:

"Nothing in recent times has caused a spike quite like that: not the Olympics (Nagano or Atlanta); not the beginning or end of the World Cup."

- Example: The Victoria's Secret Internet fashion show, May 18, 2000.

- Example: The Slashdot Effect:

  – "The spontaneous high hit rate upon a web server due to an announcement on a high volume news web site."

- Problem: Protecting other traffic on congested links.

What are the problems?
Crooks (Denial of Service Attacks):

- Example: Denial of Service attacks, February 7 and 8, 2000:

  – Attacks on a large number of major sites across the U.S.

  – "It's completely clear that the entire Internet had higher packet loss and far lower reachability for several hours." - John Quarterman.

- Problem: Limiting damage to the legitimate traffic at the site.

- Problem: Protecting the rest of the Internet.

Controlling High-Bandwidth Flows at the Congested Router

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

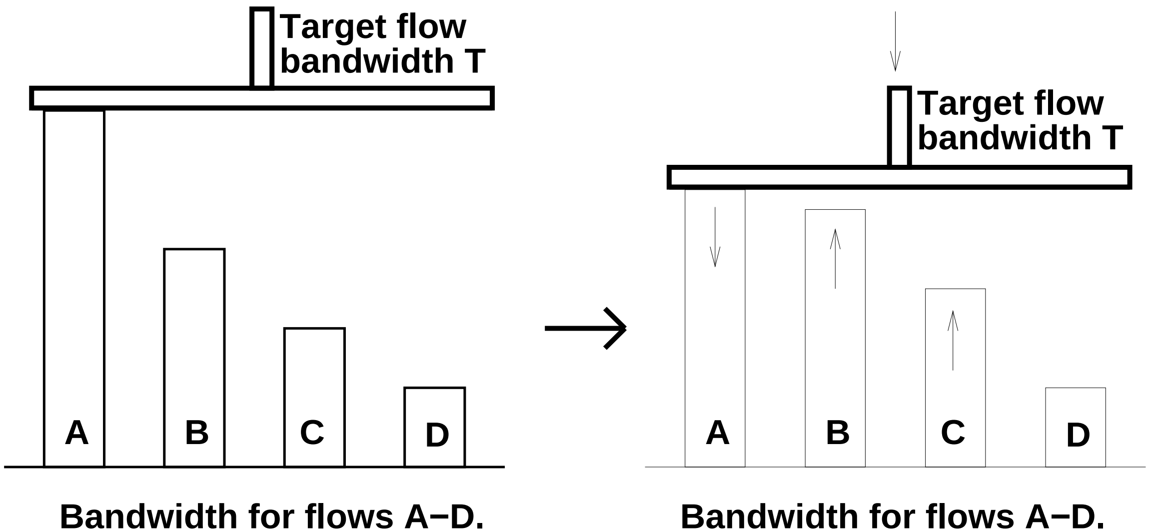

RED-PD (RED with Preferential Dropping),
restricting flows to a target bandwidth T.

Controlling High-Bandwidth Flows: Mechanisms

- Use the packet drop history at the router to detect high-bandwidth flows.

- The target bandwidth T is from the TCP throughput equation: $\frac{\sqrt{1.5}}{R\sqrt{p}}$.

  R: a configured round-trip time

  p: the current packet drop rate

.

- Monitored flows are rate-limited before the output queue.

- Monitored flows could be misbehaving flows (e.g., not using end-to-end congestion control) or conformant flows with small round-trip times.

- Identifying which monitored flows are *misbehaving* 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 rate-limiting before the output queue.

- Differences:
  - Aggregate-based congestion control (ACC) should rarely be invoked.
  - Aggregates (e.g., mobs, crooks) can have overlapping definitions.
    - E.g., dst 192.0.0.0/12, or src www.victoriasecret.com.
  - 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 flash crowd i, the drop rate is p_i (e.g., $p_i = 0.2$).
  - The throughput for A_j, for $j \neq i$, is roughly $\frac{1}{\sqrt{p_i/p}}$ of its value without the flash crowd (e.g., 1/5-th of its old value).

- Flash crowd i with ACC:
  - Assume that 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 rate-limiting 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 rate-limiting 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 by itself does not protect the other traffic in the network.

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

Open 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 probably wouldn't be invoked.

- Would the "policy knobs" in ACC be of use to ISPs?

– E.g., An aggregate could perhaps be defined as traffic to or from a neighboring ISP.