Token Bucket Regulation of VOIP Wide Area Networks

An Extension of the Paradigm Aimed at scaling issues with variable congestion points

What is QoS to the kernel developer?

- Controlling Network Congestion Points
  - Using a variety of queuing methods.
  - Assigning bandwidth to layer 2-7 entities.
  - Controlling individual packets.

- Bandwidth Control Through Queuing
  - Token Bucket Regulation supplements AltQ/CBQ.

Basic Token Bucket Regulation

Tokens regulate de-queue rate while the # of tokens per queue regulate burst.

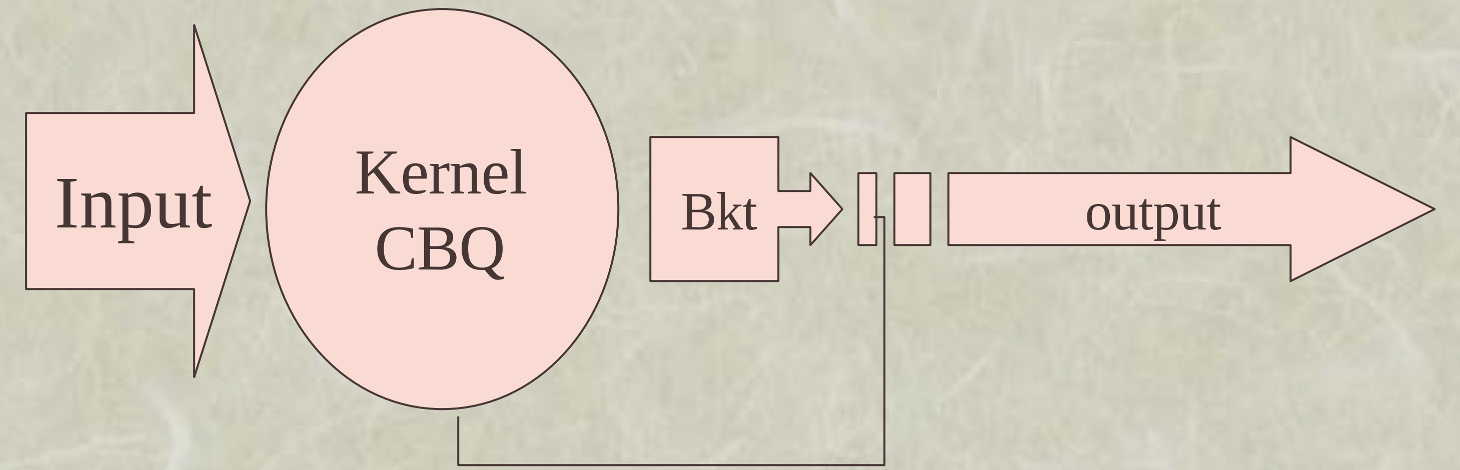

The kernel generates tokens at a set rate.
If there are available tokens then a packet is allowed out.

Token Bucket Regulated Bursting

- Bucket depth of ‘X’ Tokens
  - Allows a burst to consume X packets

Interval	1	2	3	4
Packet In	3	3	2	0
Bucket Count	5	3	1	1
Packet Out	3	3	1	1

tbrconfig Parameters

- tbrconfig *device bits/sec size*
  - Good for overall device control
  - Parameters are based on device packet rate and bit rate.
  - A device speed TBR uses transmission complete interrupts
  - Or uses a kernel generated clock
    - With rates less than device speed

TBR with Multiple Queues

- With CBQ and other QoS methodologies multiple queues can exist with different TBR parameters.
  - Underutilized bandwidth issues can be handled by bandwidth borrowing and dynamic TBR settings.
  - Netscreen uses "Double TBR" with a shared bucket of excess tokens.

TBR and VOIP 1

- VOIP and other streaming data apps are helped tremendously by TBR when coexisting with large packet, bursty apps like ftp.

- Basic TBR reduces device based congestion.
  - Particularly effective with rates < device

TBR and VOIP 2

- Multiple Queue Solutions allow for adapting size and possibly rate for a variety of bandwidth, packet size, and burst parameters.
- However VOIP and other streaming data apps have unique problems.

Streaming Data and Congestion

- Packet loss is always bad, but…
  - – Spreading loss evenly across time is potentially better than dropping established connections, refusing valid connections through conservative admission control, or dropping blocks of consecutive packets.

VOIP Information Content. 1

- Speech runs from very slow at 80 wpm to fast at 300 wpm and extremes of +700 wpm from tobacco auctioneers. Medium fast speech has a rate of 160-210 wpm.
- Phonemes per word vary depending on the complexity of the speech from 2 phonemes average in children's speech to 7 phonemes average in highly technical presentations.

VOIP Information Content. 2

- Given that wpm drops as phonemes per word rise. We can represent very rich content speech as containing :
  - (210 wpm * 5 ph/w) / 60 sec = 17.5 ph/sec
  - Assuming 30% inter-word quiet time the average phoneme length then becomes 40ms.
  - Given an average VOIP packet containing 20ms of voice the loss of a single packet is unlikely to result in the loss of a phoneme, whereas a loss of 2-3 consecutive packets likely will.

VOIP and Internet Congestion

- VOIP has particular characteristics that make it more sensitive to degradation within the cloud.
  - Delay and jitter are exaggerated due to the bi-directional nature of VOIP.
  - The controlled environment of most tests do not reflect the true nature of VOIP in the real ‘net’.

Ping: shore.net ⟺ iisc.com 1

- traceroute to shell.shore.net
- 1 dori (198.5.5.243) 2 ms 2 ms 2 ms
- 2 10.253.25.5 (10.253.25.5) 4 ms 5 ms 5 ms
- 3 12.125.47.69 (12.125.47.69) 5 ms 4 ms 4 ms
- …
- 10 cer-core-01.inet.qwest.net (205.171.205.34) 28 ms 28 ms 28 ms
- 11 205.171.139.6 (205.171.139.6) 77 ms 28 ms 28 ms
- 12 dca-core-01.inet.qwest.net (205.171.8.165) 220 ms 220 ms 220 ms
- 13 dca-edge-01.inet.qwest.net (205.171.9.22) 226 ms 219 ms 223 ms
- 14 65.122.30.142 (65.122.30.142) 218 ms 210 ms 207 ms

Ping: shore.net ⟺ iisc.com 2

- traceroute to 198.5.5.5 (198.5.5.5), 30 hops max, 40 byte packets
- 1 lynn2-br1-fa2-0-0-1.wharf.shore.net (207.244.124.10) 0.701 ms
- 2 lynn2-ar1-f1-0.wharf.shore.net (207.244.95.17) 0.772 ms 0.621 ms
- …
- 5 209.227.135.8 (209.227.135.8) 11.996 ms 11.576 ms 11.707 ms
- 6 65.122.30.141 (65.122.30.141) 11.661 ms 11.914 ms 11.721 ms
- 7 205.171.9.21 (205.171.9.21) 21.824 ms 12.113 ms 11.870 ms
- 8 205.171.9.18 (205.171.9.18) 11.668 ms 11.69 ms 11.87 ms
- 9 205.171.1.138 (205.171.1.138) 190.62 ms 190.39 ms 195.66 ms
- 10 tbr1-p013201.wswdc.ip.att.net (12.122.11.233) 198.939 ms
- …

Controlling Packet Loss

- Implemented Integrated System's method involves using a second Token Bucket Regulator to selectively drop packets.

- The 2^{nd} TBR is set to the per call determined size and rate.

TBR-2 Parameters

- TBR-2 can be set dynamically as conditions change.
- During call setup test network conditions
- Retest every ‘X’ packets
- When tested conditions meet or exceed the allocated bandwidth , then TBR-2 has no effect.
- Otherwise TBR-2 drops packets in a regular and scheduled fashion.

Packet Disposition

	QoS	TBR	TBR-2	Result
1	0	N/A	N/A	Held
2	1	0	N/A	Held
3	1	1	1	Sent
4	1	1	0	Freed

Coding TBR-2

- Coding TBR-2 largely follows TBR in terms of code placement.
- The 2 major differences are:
  - Failure de-queues to m_free
  - Dynamic setting requires userland to kernel hooks (ioctls) that are not required by TBR