RADoN: Storage Network QoS

Andrew Shewmaker[1]
Carlos Maltzahn[1]
Scott Brandt[1]

Tim Kaldewey[1]
Richard Golding[2]
Theodore M Wong[2]

[1]University of California Santa Cruz
Computer Science Department
{shewa,kalt,carlosm,scott}@cs.ucsc.edu

[1]IBM Almaden Research Center
Storage Systems Department
{rgolding,theowong}@us.ibm.com

The Goal of RADoN

Performance guarantees on standard commodity storage networks

- General
- Flexible
- Fine grained

A Canonical Storage Network

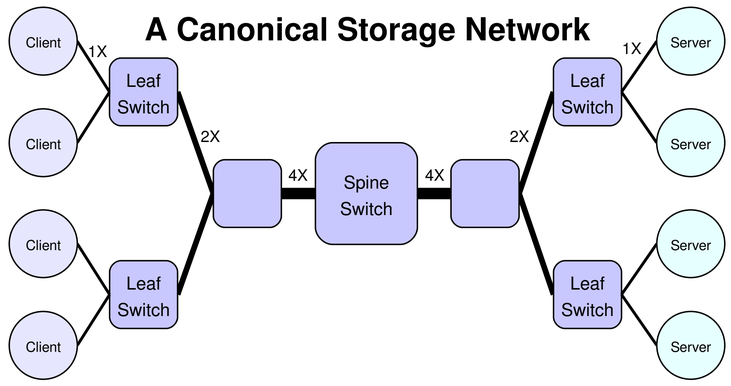

Switch Congestion is the Key Problem

Rate guarantees ensure a feasible volume of network traffic

Congestion in switches due to port contention may still cause

- large variance in delays
- packet loss

Intra and Inter Port Contention in a Simple Switch Model

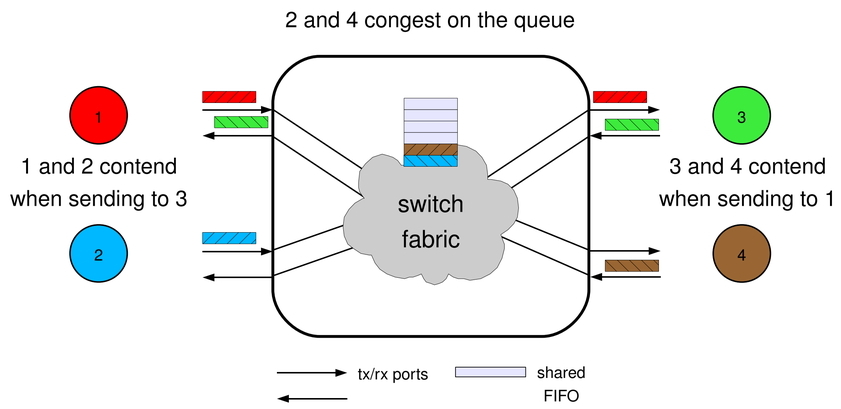

Methods of Congestion Detection

- Traditionally: packet loss, cannot prevent queue overflows
- RADoN uses *Forward Delay* as in TCP Santa Cruz

Packet Loss Based Congestion Control *vs.* Forward Delay Based Congestion Control

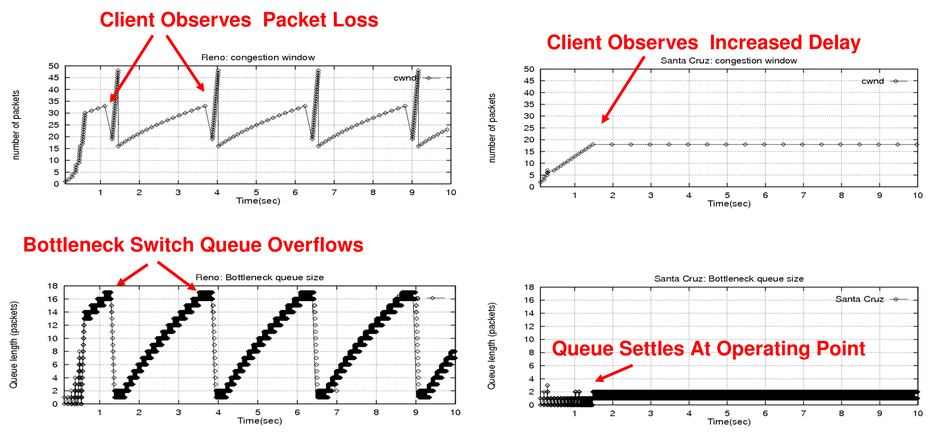

Improving TCP Congestion Control Over Internets with Heterogeneous Transmission Media (1999)
Christina Parsa, J.J. Garcia-Luna-Aceves. Proceedings of the 7th IEEE ICNP

Correct Congestion Detection and Controlled Response

- Corresponding to a share
- Bound by a time limit
- Addressing debiliating synchronization

The RAD model separates scheduling into Resource Allocation and Dispatching.
Proven correct for CPU scheduling, RADoN applies the model to the network resource.

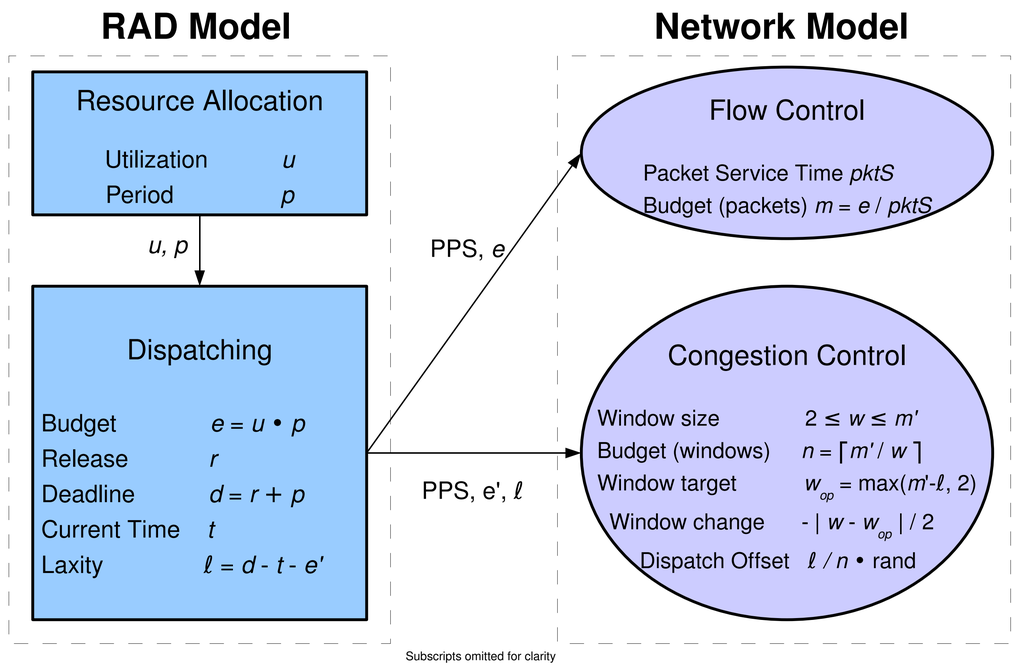

The key to RADoN is a guaranteed real-time dispatching algorithm

Least Laxity First (LLF) is a guaranteed scheduling algorithm that dispatches jobs in order of criticality.

Laxity $\ell = d - t - e'$

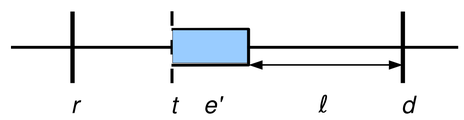

Less Laxity More (LLM) approximates LLF without requiring:

- Global knowledge
- Synchronization

A stream's work for a given period is divided into windows that change in dispatch time and size according to its laxity and remaining work when congestion is detected.

Laxity Based Congestion Response

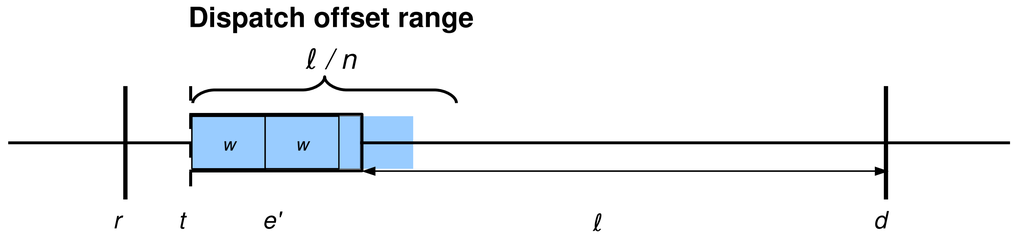

Window Size Response to Congestion

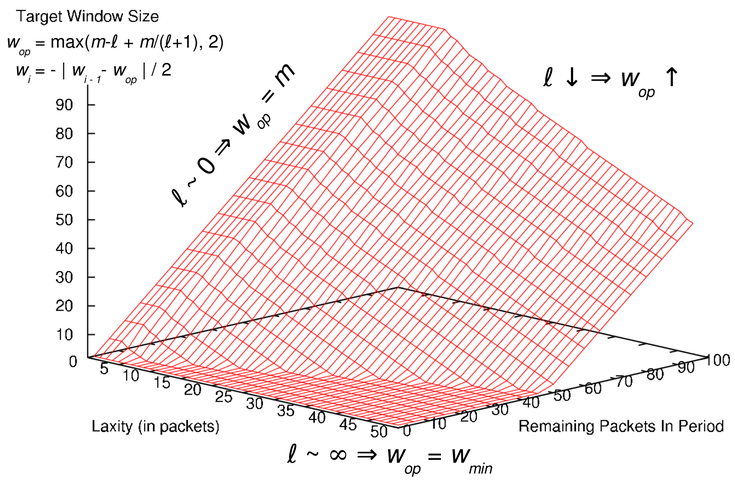

RADoN Summary

Flow = work per period

Congestion detected with *Forward Delay*

Congestion response

- controlled via window size
- bounded by laxity
- proportional to remaining work

This work was supported by Los Alamos National Laboratory and NSF Award No. CCF-0621534