


CSE398:

Network Systems Design



Instructor: Dr. Liang Cheng
Department of Computer Science and Engineering
P.C. Rossin College of Engineering & Applied Science
Lehigh University

April 20, 2005



Outline


- Recap
 - SystemC
 - APP500: classification & state engine blocks
- Traffic manager
- Summary and homework



Outline

- Recap
- **Traffic manager**
 - **Buffer management, traffic shaping, stream editing**
- Summary and homework

APP550 Overview





Primary Functions

- Buffer management
- Completion of flow policing and packet discard
- Traffic shaping
- Bandwidth allocation
- Packet modification



Buffer Management

- A packet remains in a buffer from ingress to egress
- Extra packets coming?
 - Tail drop
 - Random early detection



Random Early Detection

- Deciding the dropping probability
- Two thresholds
 - S_{\min} to start RED
 - S_{\max} to make RED non-random
 - $P = (Q_{\text{avg}} - S_{\min}) / (S_{\max} - S_{\min})$
 - Why Q_{avg} ; not Q_{current} ?
 - $Q_{\text{avg}} = aQ + (1-a)Q_{\text{avg}}$



Completion of Flow Policing

- Three steps in flow policing or traffic policing
 - Binding a packet to a flow
 - Policing engine invoked
 - Traffic manager acts based on the policing results
 - Flow profile: a set of statistical bounds of a flow
 - An out-of-profile packet is a candidate for discard
 - Different out-of-profile levels?
 - Non-preemptive discard



Traffic Profile

- Use a VBR as an example
 - Sustained bit rate (SBR)
 - Peak bit rate (PBR)
 - Sustained burst size (SBS)
 - Peak burst size (PBS)

Color	Meaning
Red	Flow exceeds the peak profile
Yellow	Flow exceeds the sustained profile, but does not exceed the peak profile
Green	Flow is below the sustained profile



WRED

- Weighted RED
 - Policer marks packets with colors
 - WRED computes the dropping probabilities

Parameter	Meaning
T_{\min}^{red}	Minimum threshold for a red packet
T_{\max}^{red}	Maximum threshold for a red packet
T_{\min}^{yellow}	Minimum threshold for a yellow packet
T_{\max}^{yellow}	Maximum threshold for a yellow packet
T_{\min}^{green}	Minimum threshold for a green packet
T_{\max}^{green}	Maximum threshold for a green packet



Implementation of WRED

- Computing the dropping probability
 - Floating point computation is costly
- $Q_{avg} = aQ + (1-a)Q_{avg}$
 - Choose a to be an inverse of a power of two
- $P = (Q_{avg} - S_{min}) / (S_{max} - S_{min})$
 - Stair-wise approximation



Stair-wise Approximation

- If $Q_{\text{avg}} \leq S_{\text{min}}$ then $P=0$
- If $S_{\text{min}} < Q_{\text{avg}} \leq S_{\text{min}} + D$ then $P=1/4$
- If $S_{\text{min}} + D < Q_{\text{avg}} \leq S_{\text{min}} + 2D$ then $p=1/2$
- If $S_{\text{min}} + 2D < Q_{\text{avg}} \leq S_{\text{min}} + 3D$ then $p=3/4$
- If $S_{\text{min}} + 3D = S_{\text{max}} < Q_{\text{avg}}$ then $p=1$




Scheduler for Traffic Shaping

- APP550 allows a programmer to create a set of scripts (a program) written in C-NP to select a packet to be transmitted from a set of queues
 - The scheduler selects the packet that will be transmitted on the next invocation
 - CBR, VBR: circular time slot table

Circular Time Slot Table

- Each entry identifies one of the queues the scheduler serves
- A global pointer indicating the next slot in the table to use
- CBR shaping
 - A scheduler sends output to an OC-12 interface
 - 6 entries in the table
 - If queue 1's CBR flow has been assigned a rate of OC-4
 - Then queue 1 should appear ? times in the table





Bandwidth Allocation

- Divide bandwidth into separate channels
- Fixed bandwidth allocation
 - Isolated channels
 - Maximum per channel is set and cannot encroach other channels even there is an idle channel
 - E.g.: TDM
- Proportional bandwidth allocation
 - Controlled sharing among channels
 - A fixed allocation scheme + excess bandwidth is divided proportionally
 - Weighted round-robin algorithm



Traffic Management Hierarchy

Level	Mechanism	Number
1	Port manager	256 total
2	Logical port	1024 total
3	Scheduler	4 per logical port
4	QoS queue	64K per scheduler
5	CoS queue	16 per QoS queue



Stream Editor Engine

- Two SED engines are last step before output.
- Except for header-trailer stripping in FPL's fQueue and fTransmit, SED is the only stage for PDU modification.
- Arguments come from DID parameter memory, fTransmit info bits, and Buffer Management & Traffic Shaper Engines.
- Output is modification of outgoing PDU.
- SED maintains no state.



Traffic Manager Data Flow

- DID selects SED script ID, parameters, queue.
- Scheduler selects Buffer Manager and Traffic Shaper scripts.
- Queue selects Buffer Manager and Traffic Shaper parameters.



Outline

- Recap
- Traffic manager
- **Summary and homework**



Reading Materials

- Examine FPL and C-NP policing code (*.asl files) in problems/learnbridge on CD.