Target-Based Power Control for Queueing Systems with Applications to Packet Switches

Benjamin Yolken, Dimitrios Tsamis, Nick Bambos
Stanford University

Introduction

Power consumption has become a major issue in modern data centers and thus power control has attracted a lot of attention.

Many components of a data center, including packet switches, can modeled as queueing systems: Jobs arrive at the system and wait for service. The system has a limited set of resources which can be deployed to service these jobs. At each point of time, the operator chooses a service configuration, i.e. a feasible combination of the device's resources.

Hence algorithms developed for packet switches can be generalized and applied to other components as well.

Goal

Design a scheduling algorithm that offers the following:

- Stability
- Differentiated Quality of Service
- Power Control

Use existing algorithms as much as possible

Notation

Q parallel queues, indexed by q = 1, 2, …, Q

X(t): vector of job backlogs in each of the queues

S(t): a service configuration (defines how many jobs are serviced from each queue)

S: set of possible service configuration

A(t): vector of arrivals at time t

$\mathbf{X}(t+1) = [\mathbf{X}(t) - \mathbf{S}(t)]^{+} + \mathbf{A}(t)$

Projective Cone Scheduling

Rule for choosing service configuration:

$\mathbf{S}(t) = \arg\max \langle \mathbf{S}, \mathbf{B}\mathbf{X}(t) \rangle$

Where **B** is a Q×Q matrix with the following properties:
- Positive Definite
- Symmetric
- Has non-positive off-diagonal elements

B = **I** gives popular MWM algorithm [3]

PCS is proven to be stable [1]

Speed Modes

Assume:

$\mathcal{S} = \{S \mid S = kS_0,\ k \in \{0,1,\dots,K\},\ S_0 \in \mathcal{S}_0\}$

S_0 is a "base" configuration with only 0/1 elements

k is the *speed mode*: as k increases more jobs are serviced from each queue

k = 0 corresponds a "shutdown" of the switch

By using a speed mode lower than the maximum one, less power is consumed, depending on the power profile of the switch.

PCS will always choose the highest k – not ideal for power control

Target-based PCS

Rule for choosing service configuration:

$\mathbf{S}(t) = \arg\max \langle \mathbf{S}, \mathbf{B}(\mathbf{X}(t) - \mathbf{b}) \rangle$

b is a vector of *backlog targets* – set according to the relative delay sensitivities of the associated job flows

Differentiated QoS is offered by playing with **b**

Limited power savings when switch goes into shutdown mode [2]

TP-PCS

Split space of possible backlog vectors in sub-spaces $X_1, \dots, X_K$, where k denotes selected speed mode for the sub-space.

Great flexibility in choice!

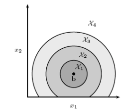

Target/Power projective cone scheduling

Set $\mathbf{S}_0(t) = \arg\max \langle \mathbf{S}, \mathbf{B}(\mathbf{X}(t) - \mathbf{b}) \rangle$
Set $k = k \mid X(t) \in X_k$
Use service vector: $S(t) = k\mathbf{S}_0(t)$

- Great flexibility in power control
- Differentiated QoS by manipulating the targets
- Stable (proof omitted)

Simulations

4×4 switch

- Assume speed mode regions X_k are cocentric spheres around the point **X** = **b**
- Use **B** = **I** (MWM) and $K = 4$ speed modes
- Power profile: speed mode k uses power k^2

Packet arrival model:
For each VOQ sample from a binomial distribution with parameters *(K, a/4)*, where $a \in \{0.3, 0.5, 0.7, 0.9\}$

Backlog targets:
βB where $\beta > 0$ and B is a random matrix whose elements are sampled uniformly from [0,1]

Simulations

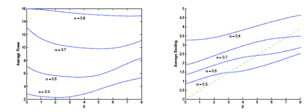

(a) Average power vs. β for fixed δ (b) Average backlog (solid) and target (dashed) vs. β for fixed δ

Fig. 3. Average power and backlog versus β (target level) for a simulated 4×4 packet switch with $K = 4$ speed modes.

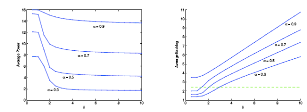

(a) Average power vs. δ for fixed β (b) Average backlog (solid) and target (dashed) vs. δ for fixed β

Bibliography

[1] K. Ross and N. Bambos "Local search scheduling algorithms for maximal throughput in packet switches", in *Proc. IEEE Infocom '04*, Mar. 2004

[2] B. Yolken and N. Bambos "Power management of packet switches via differentiated delay targets, in *Proc. IEEE ICC '08*

[3] N. McKeown, A. Mekkittikul, V. Anantharam and J. Walrand, "Achieving 100% throughput in an input-queued switch", *IEEE Transactions on Communications, Vol 47, No 8*, August 1999