Task Assignment in a Server Farm with Switching Delays and General Energy-Aware Cost Structure

Esa Hyytiä, Rhonda Righter and Samuli Aalto

Aalto University, Finland
University of California Berkeley, USA

INFORMS APS, July 2013

Aalto University
School of Electrical Engineering

Outline

Aalto University
School of Electrical Engineering

Task Assignment with Switching Delay

- Model
- Reference Policies
- First Policy Iteration Approach

Aalto University
School of Electrical Engineering

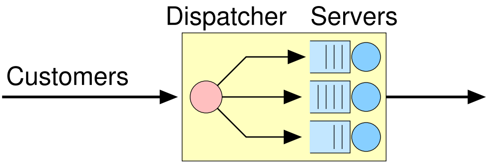

Basic problem $\mathrm{E}[\mathbf{T}]$

- k parallel queues
- Tasks arrive at rate λ (Poisson process)
- Objective: minimize latency, waiting time, ...

Aalto University
School of Electrical Engineering

Model for Server Farm

- *k* parallel servers
- Size-aware setting

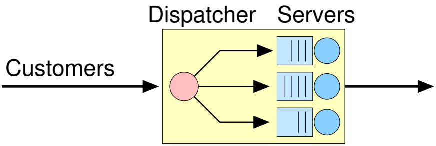

Model for Server Farm

- k parallel servers
- Size-aware setting

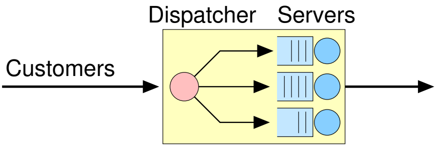

Distinctive features here

- Idle servers are switched OFF to save energy

Aalto University
School of Electrical Engineering

Switching Delays and Energy

Model for Server Farm

- k parallel servers
- Size-aware setting

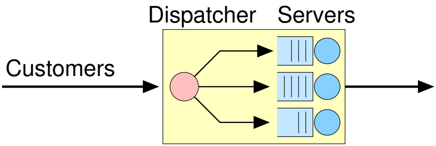

Distinctive features here

- Idle servers are switched OFF to save energy
- Switching ON delay postpones the start of the service

Aalto University
School of Electrical Engineering

Switching Delays and Energy

Model for Server Farm

- k parallel servers
- Size-aware setting

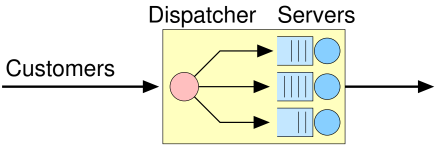

Distinctive features here

- Idle servers are switched OFF to save energy
- Switching ON delay postpones the start of the service
- Energy- and Delay-aware cost structure
  - Switching costs
  - Running costs
  - Holding costs (per job)

Aalto University
School of Electrical Engineering

Model for Server Farm

- k parallel servers
- Size-aware setting

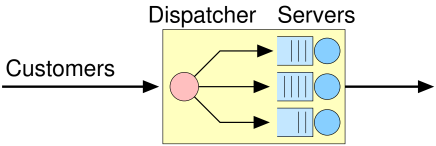

Distinctive features here

- Idle servers are switched OFF to save energy
- Switching ON delay postpones the start of the service
- Energy- and Delay-aware cost structure
  - Switching costs
  - Running costs
  - Holding costs (per job)
- Objective to balance between
  - Energy consumption
  - Performance (e.g., latency)

Aalto University
School of Electrical Engineering

Model for Server Farm

- k parallel servers
- Size-aware setting

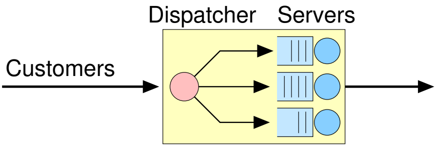

Distinctive features here

- Idle servers are switched OFF to save energy
- Switching ON delay postpones the start of the service
- Energy- and Delay-aware cost structure
  - Switching costs
  - Running costs
  - Holding costs (per job)
- Objective to balance between
  - Energy consumption
  - Performance (e.g., latency)

Heterogeneous servers, job-specific costs, ...

Aalto University
School of Electrical Engineering

Server Farms with and without Switching Delay

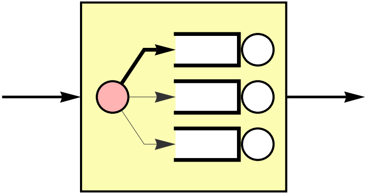

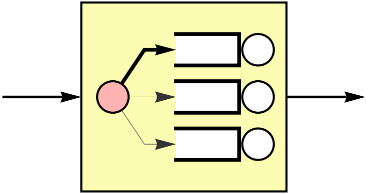

Related models (M/G/1):

- Removable servers, *N*-policy
  - (Yadin & Naor 1963; Heyman 1968)
  - Service starts when *n*th customer arrives
- Vacation models, *T*-policy
  - (Levy & Yechiali 1975; Heyman 1977)
  - Server returns periodically to check the queue
- *D*-policy, service starts when backlog exceeds *d*

Aalto University
School of Electrical
Engineering

Related work

Related models (M/G/1):

- Removable servers, N-policy
  - (Yadin & Naor 1963; Heyman 1968)
  - Service starts when nth customer arrives
- Vacation models, T-policy
  - (Levy & Yechiali 1975; Heyman 1977)
  - Server returns periodically to check the queue
- D-policy, service starts when backlog exceeds d

Results for switching delay:

- M/G/1 with setup times (Welch, Oper. Res., 1964)
- M/M/k approximations (Gandhi et al., Sigmetrics'10)
- M/M/k exact results (Gandhi et al., Sigmetrics'13)

Aalto University
School of Electrical Engineering

Related work

Related models (M/G/1):

- Removable servers, *N*-policy
  - (Yadin & Naor 1963; Heyman 1968)
  - Service starts when *n*th customer arrives
- Vacation models, *T*-policy
  - (Levy & Yechiali 1975; Heyman 1977)
  - Server returns periodically to check the queue
- *D*-policy, service starts when backlog exceeds *d*

Results for switching delay:

- M/G/1 with setup times (Welch, Oper. Res., 1964)
- M/M/*k* approximations (Gandhi et al., Sigmetrics'10)
- M/M/*k* exact results (Gandhi et al., Sigmetrics'13)

No delay- and energy-*savvy* task assignment policies!

Aalto University
School of Electrical Engineering

Definition

Static policy chooses the server independently of the queue states

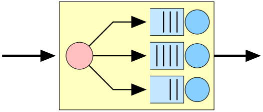

Definition

Static policy chooses the server independently of the queue states

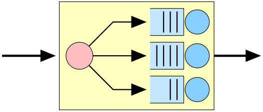

1. **Bernoulli splitting (RND)**
   Choose a queue at random using probabilities p_i

Aalto University
School of Electrical Engineering

Static Policies

Definition

Static policy chooses the server independently of the queue states

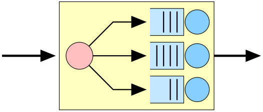

1. **Bernoulli splitting (RND)**
   Choose a queue at random using probabilities p_i
2. **Size-Interval-Task-Assignment (SITA)**
   Assignment by the **queue-specific** ranges of job sizes.

   "Short jobs to Queue 1 and the rest to Queue 2"

   - Proposed in Crovella et. al (Sigmetrics'98) and Harchol-Balter et. al (J. of PDC, 1999)
   - Optimal size-aware state-free for FCFS (Feng et. al, -05)

Aalto University
School of Electrical Engineering

Dynamic Policies

Definition

Actions of a **dynamic policy** depend on the queue states

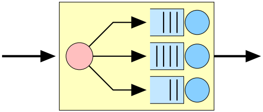

Aalto University
School of Electrical Engineering

Dynamic Policies

Definition

Actions of a **dynamic policy** depend on the queue states

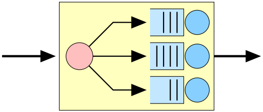

1. **Join-the-Shortest-Queue (JSQ)**
   Optimal when Poisson arrivals, Exp-distributed job sizes, identical servers, and the queue occupancy is known
   (Haight 1958; Winston 1977)

Aalto University
School of Electrical Engineering

Definition

Actions of a **dynamic policy** depend on the queue states

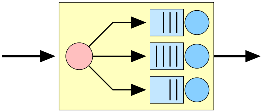

1. **Join-the-Shortest-Queue (JSQ)**
   Optimal when Poisson arrivals, Exp-distributed job sizes, identical servers, and the queue occupancy is known
   (Haight 1958; Winston 1977)

2. **Round-robin (RR)**
   Optimal with identical servers initially in the same state, known routing history and unknown queue occupancy
   (Ephremides et. al, 1980; Liu&Towsley-94; Liu&Righter -98)

Aalto University
School of Electrical Engineering

Definition

Actions of a **dynamic policy** depend on the queue states

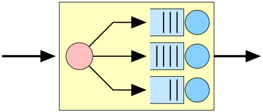

1. **Join-the-Shortest-Queue (JSQ)**
   Optimal when Poisson arrivals, Exp-distributed job sizes, identical servers, and the queue occupancy is known
   (Haight 1958; Winston 1977)
2. **Round-robin (RR)**
   Optimal with identical servers initially in the same state, known routing history and unknown queue occupancy
   (Ephremides et. al, 1980; Liu&Towsley-94; Liu&Righter -98)
3. **Least-Work-Left (LWL)**
   Pick the queue with the shortest backlog
   (Sharifnia 1997)

Aalto University
School of Electrical Engineering

First Policy Iteration (FPI)

- **Decomposition** with a **static basic policy** α
  - Queues receive jobs according to Poisson processes

First Policy Iteration (FPI)

- **Decomposition** with a **static basic policy** α
  - Queues receive jobs according to Poisson processes
  - Value function is a sum of the queue-specific value functions

$$v_{\mathbf{z}}(\alpha) = \sum_i v_{\mathbf{z}_i}^{(i)}(\alpha)$$

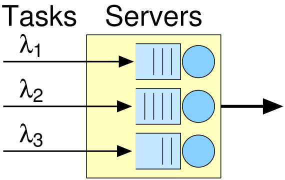

First Policy Iteration (FPI)

- **Decomposition** with a **static basic policy** α
  - Queues receive jobs according to Poisson processes
  - Value function is a sum of the queue-specific value functions

$$v_{\mathbf{z}}(\alpha) = \sum_i v_{\mathbf{z}_i}^{(i)}(\alpha)$$

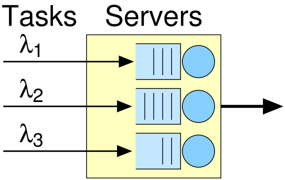

- **FPI** gives a new policy α'

$$\alpha'(x) = \underset{i}{\text{argmin}}\ \ h_i(x) + \left(v_{\mathbf{z}_i \oplus x}^{(i)} - v_{\mathbf{z}_i}^{(i)}\right)$$

where

- $h_i(x)$ is the *immediate cost* of choosing Queue i for Job x
- $v_{\mathbf{z}_i \oplus x}^{(i)} - v_{\mathbf{z}_i}^{(i)}$ is the mean increase in future costs in Queue i

Aalto University
School of Electrical Engineering

First Policy Iteration (FPI)

- **Decomposition** with a **static basic policy** $\boldsymbol{\alpha}$
  - Queues receive jobs according to Poisson processes
  - Value function is a sum of the queue-specific value functions

$$v_{\mathbf{z}}(\alpha) = \sum_i v_{\mathbf{z}_i}^{(i)}(\alpha)$$

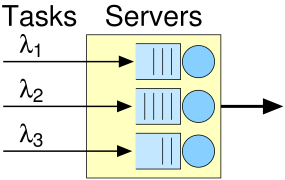

- **FPI** gives a new policy α'

$$\alpha'(x) = \underset{i}{\text{argmin}} \;\; h_i(x) + \left(v_{\mathbf{z}_i \oplus x}^{(i)} - v_{\mathbf{z}_i}^{(i)} \right)$$

where

- $h_i(x)$ is the *immediate cost* of choosing Queue i for Job x
- $v_{\mathbf{z}_i \oplus x}^{(i)} - v_{\mathbf{z}_i}^{(i)}$ is the mean increase in future costs in Queue i

Idea: $\boxed{\text{The new dynamic policy } \alpha' \text{ is better than } \alpha}$

Aalto University
School of Electrical Engineering

FPI Approach

Queueing systems

M/M/s	Krishnan, CDC (1987)
M/M/1	Aalto&Virtamo, NTS-13 (1996)
M/G/1-FCFS	Sassen et al., Neerlandica (1997)
M/M/1 & M/M/1/N	Koole, CDC (1998)
M/Cox(r)/1	Bhulai, JAP (2006)

Aalto University
School of Electrical
Engineering

FPI Approach

Queueing systems

M/M/s	Krishnan, CDC (1987)
M/M/1	Aalto&Virtamo, NTS-13 (1996)
M/G/1-FCFS	Sassen et al., Neerlandica (1997)
M/M/1 & M/M/1/N	Koole, CDC (1998)
M/Cox(r)/1	Bhulai, JAP (2006)

Size-aware queueing systems

M/G/1 FCFS/LCFS/SRPT	Hyytiä et al., EJOR (2012), Sigmetrics (2012)
M/G/1 (wrt. energy)	Penttinen et al., IPCCC (2011)
M/D/1-PS	Hyytiä et al., ITC (2011)
M/M/1-PS	Hyytiä et al., Performance (2011)

Aalto University
School of Electrical Engineering

FPI Approach

Queueing systems

M/M/s	Krishnan, CDC (1987)
M/M/1	Aalto&Virtamo, NTS-13 (1996)
M/G/1-FCFS	Sassen et al., Neerlandica (1997)
M/M/1 & M/M/1/N	Koole, CDC (1998)
M/Cox(r)/1	Bhulai, JAP (2006)

Size-aware queueing systems

M/G/1 FCFS/LCFS/SRPT	Hyytiä et al., EJOR (2012), Sigmetrics (2012)
M/G/1 (wrt. energy)	Penttinen et al., IPCCC (2011)
M/D/1-PS	Hyytiä et al., ITC (2011)
M/M/1-PS	Hyytiä et al., Performance (2011)

Blocking systems

M/M/s/s	Krishnan, CDC (1986)
M/M/s/k	Leeuwaarden et al. (2001)

Aalto University
School of Electrical
Engineering

Analysis of Single M/G/1

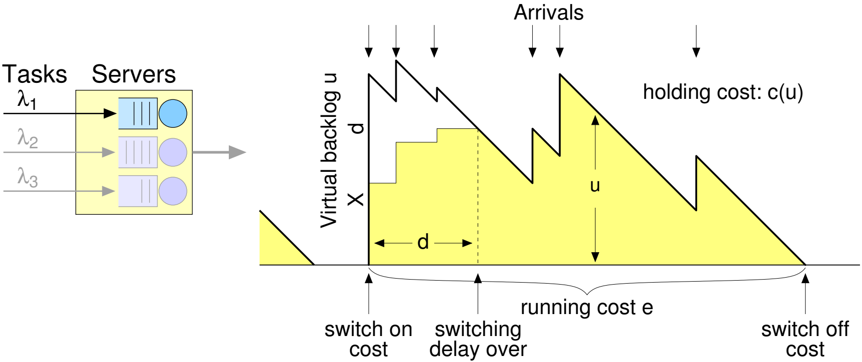

[1] See (Heyman 1968) and (Feinberg & Kella, 2002) for M/G/1.

Cost Structure

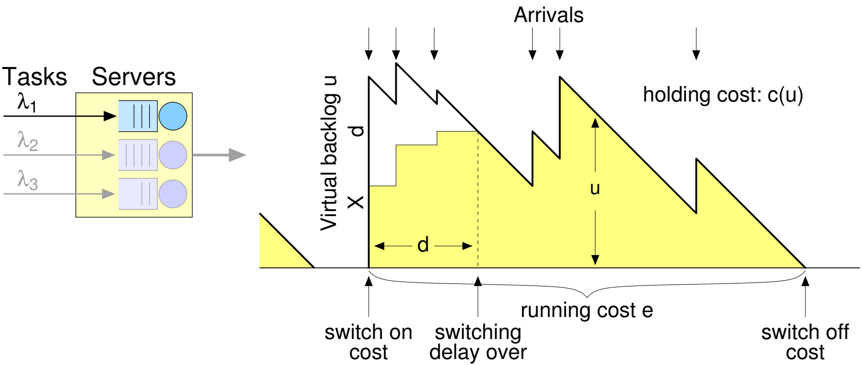

The queue-specific cost structure[1]

(i) **switching costs** $(k_{\mathrm{on}}, k_{\mathrm{off}})$ (per cycle)

[1] See (Heyman 1968) and (Feinberg & Kella, 2002) for M/G/1.

Cost Structure

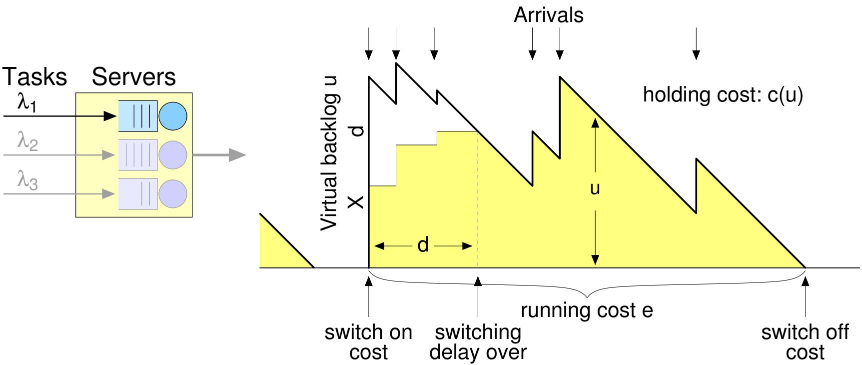

The queue-specific cost structure[1]

(i) **switching costs** ($k_{\mathrm{on}}, k_{\mathrm{off}}$) (per cycle)

(ii) **running costs** ($e_{\mathrm{on}}, e_{\mathrm{off}}$) (per unit time)

[1] See (Heyman 1968) and (Feinberg & Kella, 2002) for M/G/1.

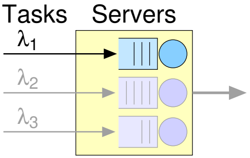

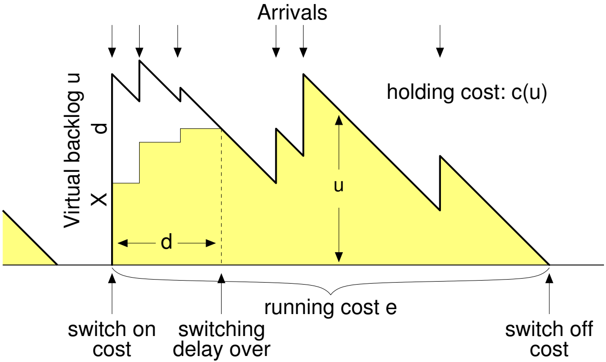

The queue-specific cost structure[1]

(i) **switching costs** $(k_{\text{on}}, k_{\text{off}})$ (per cycle)

(ii) **running costs** $(e_{\text{on}}, e_{\text{off}})$ (per unit time)

(iii) **holding cost** $c(u)$ (per unit time), u = virtual backlog

[1] See (Heyman 1968) and (Feinberg & Kella, 2002) for M/G/1.

Aalto University
School of Electrical Engineering

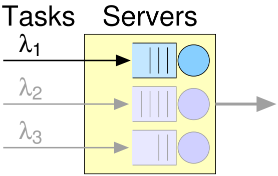

- Formally

$$v_{\mathbf{z}} \triangleq \lim_{t\to\infty} \mathrm{E}[V_{\mathbf{z}}(t) - r \cdot t]$$

where

  - $V_{\mathbf{z}}(t)$ = costs incurred during $(0, t)$ when initially in state $\mathbf{z}$
  - r = long-run mean cost rate

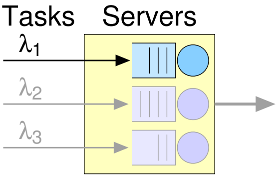

- Formally

$$v_{\mathbf{z}} \triangleq \lim_{t\to\infty} \mathrm{E}[V_{\mathbf{z}}(t) - r \cdot t]$$

  where
  - $V_{\mathbf{z}}(t)$ = costs incurred during $(0, t)$ when initially in state $\mathbf{z}$
  - r = long-run mean cost rate

- With FCFS, a sufficient state description is (u, e)
  - u = virtual backlog (measured in time)
  - $e = \begin{cases} 1, & \text{if the server is available} \\ 0, & \text{otherwise (on vacation)} \end{cases}$

Aalto University
School of Electrical
Engineering

Cost	Mean rate $\mathbf{r}_*$	Value function $\mathbf{v}_*(\mathbf{u}) - \mathbf{v}_*(\mathbf{0})$
Switching	$\lambda(1-\rho) \cdot k$	$-\lambda u \cdot k$
Running	$\rho \cdot e$	$u \cdot e$
Holding H_1	$\dfrac{\lambda\,\mathrm{E}[X^2]}{2(1-\rho)}$	$\dfrac{u^2}{2(1-\rho)}$

Aalto University
School of Electrical Engineering

Results for M/G/1 with Switching Delay

Cost	Mean rate $\mathbf{r}_*$	Value function $\mathbf{v}_*(\mathbf{u}) - \mathbf{v}_*(\mathbf{0})$
Switching	$\dfrac{\lambda(1-\rho)}{1+\lambda d} \cdot k$	$-\dfrac{\lambda u}{1+\lambda d} \cdot k$
Running	$\dfrac{\rho+\lambda d}{1+\lambda d} \cdot e$	$\dfrac{u}{1+\lambda d} \cdot e$
Holding H_1	$\dfrac{\lambda\,\mathrm{E}[X^2]}{2(1-\rho)} + \dfrac{d(2\rho+\lambda d)}{2(1+\lambda d)}$	$\dfrac{u^2}{2(1-\rho)} - \dfrac{d(2\rho+\lambda d)\cdot u}{2(1-\rho)(1+\lambda d)}$

Aalto University
School of Electrical Engineering

Results for M/G/1 with Switching Delay

Cost	Mean rate $\mathbf{r}_*$	Value function $\mathbf{v}_*(\mathbf{u}) - \mathbf{v}_*(\mathbf{0})$
Switching	$\dfrac{\lambda(1-\rho)}{1+\lambda d} \cdot k$	$-\dfrac{\lambda u}{1+\lambda d} \cdot k$
Running	$\dfrac{\rho+\lambda d}{1+\lambda d} \cdot e$	$\dfrac{u}{1+\lambda d} \cdot e$
Holding H_1	$\dfrac{\lambda\,\mathrm{E}[X^2]}{2(1-\rho)} + \dfrac{d(2\rho+\lambda d)}{2(1+\lambda d)}$	$\dfrac{u^2}{2(1-\rho)} - \dfrac{d(2\rho+\lambda d)\cdot u}{2(1-\rho)(1+\lambda d)}$

Note

- Holding cost with $d = 0$ is the Pollazcek-Khinchine formula
- Switching delay shows up as an extra term in r_{H1} and $v_{H1}(u)$
  - Extra cost in $v_{H1}(u)$ due to switching delay $\propto u$
- Decomposition property (Fuhrmann & Cooper, 1985)

- Linear holding cost corresponds to metrics such as
  - Latency (i.e., delay, sojourn time, waiting time)
  - Slowdown (ratio of the latency to job size, T/X)
  - ...anything that is directly proportional to T

Aalto University
School of Electrical Engineering

Quadratic Holding Costs

- Linear holding cost corresponds to metrics such as
  - Latency (i.e., delay, sojourn time, waiting time)
  - Slowdown (ratio of the latency to job size, T/X)
  - ...anything that is directly proportional to T
- Not everything is linear
  - E.g., longer waiting may cause more customer dissatisfaction $\Rightarrow$ cost rate increases!

Aalto University
School of Electrical Engineering

Quadratic Holding Costs

- Linear holding cost corresponds to metrics such as
  - Latency (i.e., delay, sojourn time, waiting time)
  - Slowdown (ratio of the latency to job size, T/X)
  - ...anything that is directly proportional to T
- Not everything is linear
  - E.g., longer waiting may cause more customer dissatisfaction $\Rightarrow$ cost rate increases!
- What about quadratic costs?

Virtual backlog,	cost rate $\propto U(t)^2$
Latency of Job i,	cost incurred $\propto (T_i)^2$

Aalto University
School of Electrical Engineering

Quadratic Holding Costs

- Linear holding cost corresponds to metrics such as
  - Latency (i.e., delay, sojourn time, waiting time)
  - Slowdown (ratio of the latency to job size, T/X)
  - ...anything that is directly proportional to T
- Not everything is linear
  - E.g., longer waiting may cause more customer dissatisfaction $\Rightarrow$ cost rate increases!
- What about quadratic costs?

Virtual backlog,	cost rate $\propto U(t)^2$
Latency of Job i,	cost incurred $\propto (T_i)^2$

Good news: These can be computed too!

Aalto University
School of Electrical Engineering

The mean holding cost rate is

$$
\begin{aligned}
r_{H2} &= \mathrm{E}[U^2] \\
&= \frac{3\lambda^2\,\mathrm{E}[X^2]^2 + 2\lambda(1-\rho)\,\mathrm{E}[X^3]}{6(1-\rho)^2} + \underbrace{\frac{3\rho+\lambda d}{3(1+\lambda d)}d^2 + \frac{\lambda(2+\lambda d)\,\mathrm{E}[X^2]}{2(1-\rho)(1+\lambda d)}d}_{\text{switching delay}}
\end{aligned}
$$

Aalto University
School of Electrical Engineering

Quadratic Holding Costs

The mean holding cost rate is

$$
\begin{aligned}
r_{H2} &= \mathrm{E}[U^2] \\
&= \frac{3\lambda^2\,\mathrm{E}[X^2]^2+2\lambda(1-\rho)\,\mathrm{E}[X^3]}{6(1-\rho)^2} + \underbrace{\frac{3\rho+\lambda d}{3(1+\lambda d)}d^2 + \frac{\lambda(2+\lambda d)\,\mathrm{E}[X^2]}{2(1-\rho)(1+\lambda d)}d}_{\text{switching delay}}
\end{aligned}
$$

The corresponding value function is

$$
v_{H2}(u) - v_{H2}(0) =
$$

$$
\frac{1}{3(1-\rho)}u^3+\frac{\lambda\,\mathrm{E}[X^2]}{2(1-\rho)^2}u^2-\underbrace{\left(\frac{3\rho+\lambda d}{3(1-\rho)(1+\lambda d)}d^2+\frac{\lambda(2+\lambda d)\,\mathrm{E}[X^2]}{2(1-\rho)^2(1+\lambda d)}d\right)u}_{\text{switching delay}}
$$

Aalto University
School of Electrical Engineering

Quadratic Holding Costs

The mean holding cost rate is

$$
\begin{aligned}
r_{H2} &= \mathrm{E}[U^2] \\
&= \frac{3\lambda^2\,\mathrm{E}[X^2]^2 + 2\lambda(1-\rho)\,\mathrm{E}[X^3]}{6(1-\rho)^2} + \underbrace{\frac{3\rho+\lambda d}{3(1+\lambda d)} d^2 + \frac{\lambda(2+\lambda d)\,\mathrm{E}[X^2]}{2(1-\rho)(1+\lambda d)} d}_{\text{switching delay}}
\end{aligned}
$$

The corresponding value function is

$$
v_{H2}(u) - v_{H2}(0) =
$$

$$
\frac{1}{3(1-\rho)} u^3 + \frac{\lambda\,\mathrm{E}[X^2]}{2(1-\rho)^2} u^2 - \underbrace{\left(\frac{3\rho+\lambda d}{3(1-\rho)(1+\lambda d)} d^2 + \frac{\lambda(2+\lambda d)\,\mathrm{E}[X^2]}{2(1-\rho)^2(1+\lambda d)} d \right) u}_{\text{switching delay}}
$$

- Mean cost rate (cf. PK) and value function resemble each other
- Switching delay appears as extra terms in both
- In value function, cost of switching delay proportional to $-u$

Aalto University
School of Electrical Engineering

- For an arbitrary cost function $c(u)$

$$c_1 \triangleq \mathrm{E}[c(W_1) + \ldots + c(W_{N_u})]$$

$$c_2 \triangleq \lambda\,\mathrm{E}[\int_0^{B_u} c(U_t)\,dt]$$

PASTA $\Rightarrow c_1 = c_2$

Aalto University
School of Electrical Engineering

- For an arbitrary cost function $c(u)$

$$c_1 \triangleq \mathrm{E}[c(W_1) + \ldots + c(W_{N_u})]$$

$$c_2 \triangleq \lambda \, \mathrm{E}[\int_0^{B_u} c(U_t)\, dt]$$

PASTA $\Rightarrow c_1 = c_2$

- For waiting time W and its square

Linear	$v_W(u) - v_W(0) = \lambda \left(v_{H1}(u) - v_{H1}(0) - \frac{du}{1+\lambda d} \right)$
Quadratic	$v_{W2}(u) - v_{W2}(0) = \lambda \left(v_{H2}(u) - v_{H2}(0) - \frac{d^2 u}{1+\lambda d} \right)$

Aalto University
School of Electrical Engineering

- For an arbitrary cost function $c(u)$

$$c_1 \triangleq \mathrm{E}[c(W_1) + \ldots + c(W_{N_u})]$$

$$c_2 \triangleq \lambda\,\mathrm{E}[\int_0^{B_u} c(U_t)\,dt]$$

PASTA $\Rightarrow c_1 = c_2$

- For waiting time W and its square

Linear	$v_W(u) - v_W(0) = \lambda\left(v_{H1}(u) - v_{H1}(0) - \frac{du}{1+\lambda d}\right)$
Quadratic	$v_{W2}(u) - v_{W2}(0) = \lambda\left(v_{H2}(u) - v_{H2}(0) - \frac{d^2u}{1+\lambda d}\right)$

- For latency, $v_T(u) - v_T(0) = v_W(u) - v_W(0)$
  Similarly, an expression for $v_{T2}(u)$ can be obtained

Aalto University
School of Electrical
Engineering

So what do we have?

Cost type	mean rate	value function	immediate cost
Switching cost	✓	✓	✓
Running cost	✓	✓	
Waiting time W	✓	✓	✓
Waiting time W^2	✓	✓	✓
Latency T	✓	✓	✓
Latency T^2	✓	✓	✓

$\Rightarrow$ $\boxed{\text{FPI-policy based on a general cost structure!}}$

Aalto University
School of Electrical Engineering

Solving the Task Assignment Problem

Numerical Examples

Example homogeneous server system

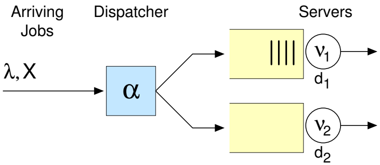

Two homogeneous servers

	service rate	switching delay
Queue 1:	$\nu_1 = 1$,	$d_1 = 1$
Queue 2:	$\nu_2 = 1$,	$d_2 = 1$

Aalto University
School of Electrical Engineering

Example homogeneous server system

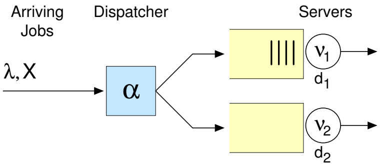

Two homogeneous servers

	service rate	switching delay
Queue 1:	$\nu_1 = 1$,	$d_1 = 1$
Queue 2:	$\nu_2 = 1$,	$d_2 = 1$

- $\lambda = 1.5$ and $\mathrm{E}[X] = 1$
- Minimize waiting time W

Aalto University
School of Electrical Engineering

Example homogeneous server system

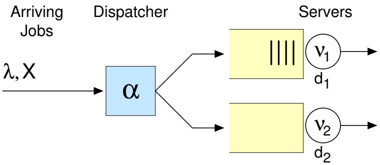

Two homogeneous servers

	service rate	switching delay
Queue 1:	$\nu_1 = 1$,	$d_1 = 1$
Queue 2:	$\nu_2 = 1$,	$d_2 = 1$

- $\lambda = 1.5$ and $\mathrm{E}[X] = 1$
- Minimize waiting time W
- Basic policy α = RND

Aalto University
School of Electrical Engineering

Example homogeneous server system

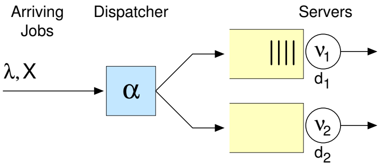

Two homogeneous servers

	service rate	switching delay
Queue 1:	$\nu_1 = 1$,	$d_1 = 1$
Queue 2:	$\nu_2 = 1$,	$d_2 = 1$

- $\lambda = 1.5$ and $\mathrm{E}[X] = 1$
- Minimize waiting time W
- Basic policy α = RND
- Server 1 busy, $u_1 > 0$
  Server 2 idle, $u_2 = 0$

Example homogeneous server system

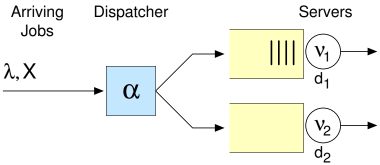

Two homogeneous servers

	service rate	switching delay
Queue 1:	$\nu_1 = 1,$	$d_1 = 1$
Queue 2:	$\nu_2 = 1,$	$d_2 = 1$

- $\lambda = 1.5$ and $\mathrm{E}[X] = 1$
- Minimize waiting time W
- Basic policy α = RND
- Server 1 busy, $u_1 > 0$
  Server 2 idle, $u_2 = 0$
- FPI sends a job to idle
  Server 2 earlier than LWL

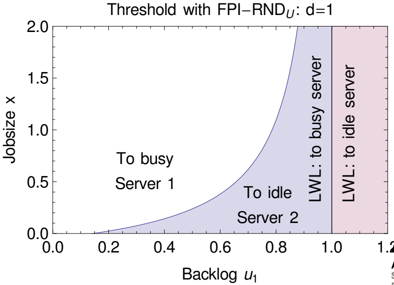

Aalto University
School of Electrical Engineering

Example heterogeneous server system

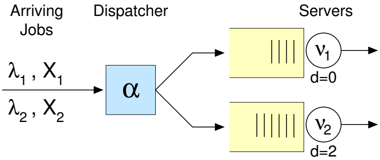

Two traffic classes:

	arrival rate	job size
Class 1:	λ_1=0.8 (fixed),	$\mathrm{E}[X_1]$=1
Class 2:	λ_2=0 . . . 0.5,	$\mathrm{E}[X_2]$=2

Two heterogeneous servers:

	service rate	switching delay
Queue 1:	$\nu_1 = 1$,	$d_1 = 0$ (none)
Queue 2:	$\nu_2 = 1$,	$d_2 = 2$

Aalto University
School of Electrical Engineering

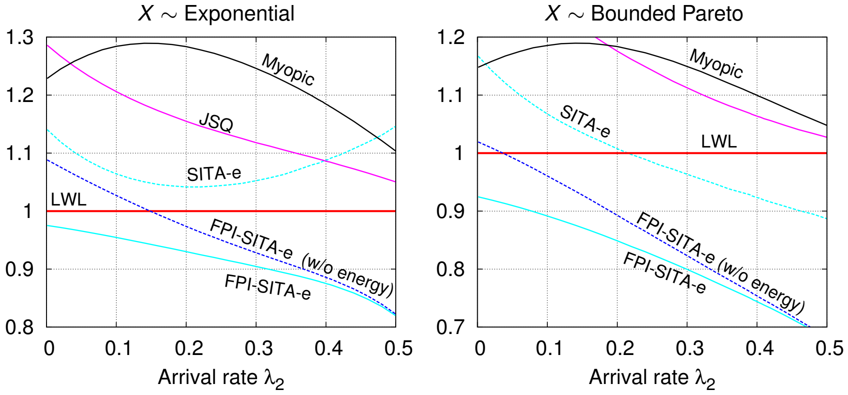

Figure: Mean cost rate with the joint-objective of **waiting time** and **running costs**: $e_1 = 1$ and $e_2 = 2$

Conclusions

- M/G/1 queue with switching delay analyzed

Aalto University
School of Electrical Engineering

Conclusions

- M/G/1 queue with switching delay analyzed
- Value functions derived with respect to
  1. Switching costs [1/cycle]
  2. Running costs [1/time]
  3. Virtual backlog U_t yielding
     - Waiting time W and its square W^2
     - Latency T and its square T^2

Aalto University
School of Electrical Engineering

Conclusions

- M/G/1 queue with switching delay analyzed
- Value functions derived with respect to
  1. Switching costs [1/cycle]
  2. Running costs [1/time]
  3. Virtual backlog U_t yielding
     - Waiting time W and its square W^2
     - Latency T and its square T^2
- Enables efficient **task assignment** taking into account
  - Switching delays and service rates
  - Current state of the system
  - Job- and server-specific cost parameters
  - Anticipated future arrivals

Aalto University
School of Electrical Engineering

Conclusions

- M/G/1 queue with switching delay analyzed
- Value functions derived with respect to
  1. Switching costs [1/cycle]
  2. Running costs [1/time]
  3. Virtual backlog U_t yielding
     - Waiting time W and its square W^2
     - Latency T and its square T^2
- Enables efficient **task assignment** taking into account
  - Switching delays and service rates
  - Current state of the system
  - Job- and server-specific cost parameters
  - Anticipated future arrivals
- Future work: active control of idle servers

Aalto University
School of Electrical Engineering

Conclusions

- M/G/1 queue with switching delay analyzed
- Value functions derived with respect to
  1. Switching costs [1/cycle]
  2. Running costs [1/time]
  3. Virtual backlog U_t yielding
     - Waiting time W and its square W^2
     - Latency T and its square T^2
- Enables efficient **task assignment** taking into account
  - Switching delays and service rates
  - Current state of the system
  - Job- and server-specific cost parameters
  - Anticipated future arrivals
- Future work: active control of idle servers

Thank you!

Aalto University
School of Electrical Engineering

References

1. Hyytiä, Righter, Aalto, *"Task Assignment in a Server Farm with Switching Delays and General Energy-Aware Cost Structure"*, submitted, 2013.
2. Hyytiä, Aalto, Penttinen, *"Minimizing Slowdown in Heterogeneous Size-Aware Dispatching Systems"*, ACM SIGMETRICS'12.
3. Hyytiä, Penttinen, Aalto, *"Size- and State-Aware Dispatching Problem with Queue-Specific Job Sizes"*, EJOR 217(2), 357–370, 2012.

See also,

`http://www.netlab.hut.fi/~esa/dispatching.html`

Aalto University
School of Electrical Engineering