Aalto University
School of Electrical
Engineering

Round-Robin Routing Policy:
Value Functions and Mean Performance with Job- and Server-specific Costs

Esa Hyytiä & Samuli Aalto
Aalto University, Finland

ValueTools 2013
10–12 December 2013
Turin, Italy

Part I
Dispatching problem

Dispatching problem

- Dispatching = Task assignment = Routing
  - *m* parallel servers with their own queues
  - random job arrivals with random service requirements
  - dispatching decisions made upon the arrival time
  - minimize e.g. mean waiting/sojourn time or mean slowdown
  - ICT applications: web server farms, supercomputer grids, etc.

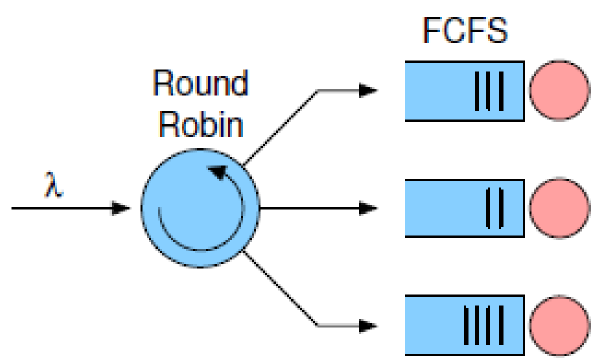

Round-Robin (RR) routing

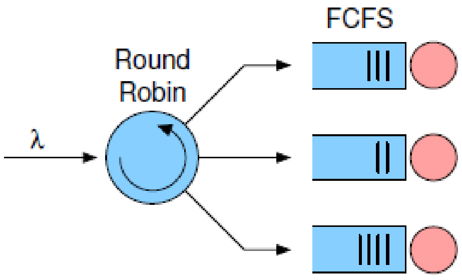

- RR assigns jobs to queues in the sequential order (1, 2, …, m, 1,2, …, m, 1, 2, …)
  - better than pure random dispatching when the servers are identical
  - optimal for identical servers when the service requirements are deterministic
    - in this case, RR is equal to JSQ and LWL
- With Poisson arrivals,
  - each queue is an Erl(m,λ)/G/1 queue (but not independent)
- Our target:
  - to improve RR by utilizing size, cost and state information

Cost structure

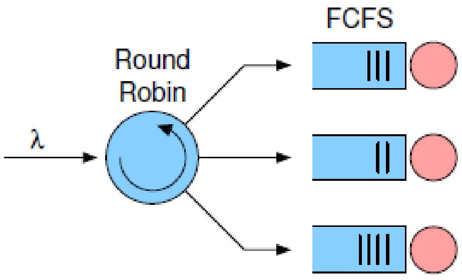

- If job j is assigned to queue k, then
  - service fee of S_{jk} is paid (once)
  - holding costs are incurred with rate H_{jk} during the waiting time
  - service time X_{jk} may be queue-specific
- Vector triplets $(\mathbf{X}_j, \mathbf{H}_j, \mathbf{S}_j)$ i.i.d.
  - but the components may depend on each other
- Examples:
  - If $H_{jk} = 1$ & $S_{jk} = 0$, then the mean waiting time minimized
  - If $H_{jk} = 1$ & $S_{jk} = X_{jk}$, then the mean sojourn time minimized
  - If $H_{jk} = 1/X_{jk}$ & $S_{jk} = 1$, then the mean slowdown minimized

First Policy Iteration (FPI)

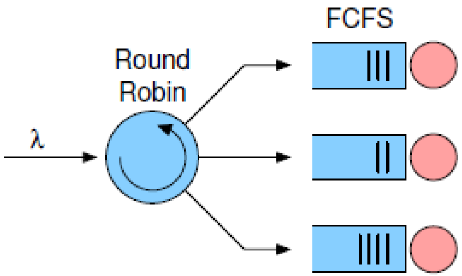

- Assume:
  - Poisson arrivals, RR routing (as the basic policy to be improved) and FCFS scheduling (locally in each queue)
- First Policy Iteration (FPI) based on the MDP theory:
  - Determine the size-aware relative values for the parallel Erl(m,λ)/G/1 queues
  - Evaluate the decision (to dispatch an arriving job to a queue) by utilizing these relative values
  - Dispatch the arriving job to the queue that minimizes the mean additional costs
- FPI-RR is a state-dependent dispatching policy
  - performs better than the original (state-independent) RR policy

Part II
Erl(m,λ)/G/1 analysis: Value functions

Value function related to service costs (1)

- Let i denote the current phase of the arrival process
- Definition 1:
  The value function v_i gives the expected difference in the infinite horizon cumulative service costs between
  - the system initially in state i and
  - the system initially in equilibrium,

$$v_i = \lim_{t\to\infty} E[V_i(t) - r_s t]$$

- Here r_s denotes the average service cost rate:

$$r_s = \frac{\lambda E[S]}{m}$$

Value function related to service costs (2)

- Assume:
  - Arrivals at the end of the final phase *m*
- Proposition 1:

$$v_i = \frac{2i - m - 1}{2m} E[S]$$

- Proof (ideas):

$$v_i = \frac{m-i+1}{\lambda}(0 - r_s) + E[S] + v_1$$

$$\frac{1}{m}(v_1 + \ldots + v_m) = 0$$

Value function related to virtual waiting costs (1)

- Let u denote the current backlog of the queue
  - backlog = virtual waiting time = unfinished work (in time units)
- Definition 2:
  The value function $v_i(u)$ gives the expected difference in the infinite horizon cumulative virtual waiting costs btw
  - the system initially in state (i,u) and
  - the system initially in equilibrium,

$$v_i(u) = \lim_{t\to\infty} E[V_i(u,t) - r_b t]$$

- Here r_b denotes the average virtual waiting cost rate

Value function related to virtual waiting costs (2)

- Remark:
  - Virtual waiting cost rate is equal to the backlog whenever the arrival process is in the final phase *m* (otherwise cost rate is 0)

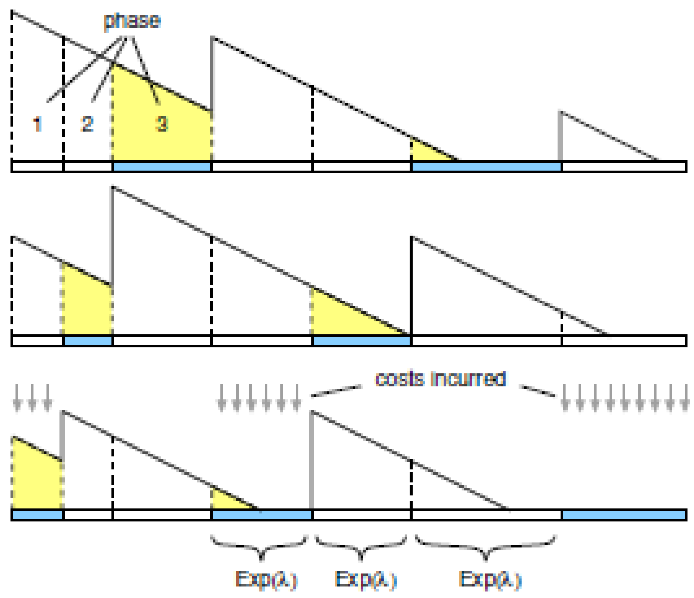

Value function related to virtual waiting costs (3)

- Proposition 2:

$$v_i'(u) = -r_b + \lambda(v_{i+1}(u) - v_i(u)), \quad i = 1,\ldots,m-1$$

$$v_m'(u) = u - r_b + \lambda \int_0^\infty (v_1(u+x) - v_m(u))\,dF(x)$$

- Remarks:
  - An efficient numerical method (based on Prop. 2) to determine the relative values $v_i(u) - v_1(0)$ and the average cost rate r_b given in the paper
  - As a useful spinoff, the mean waiting time $E[W] = m\, r_b$ in the original parallel queueing system becomes determined

Part III
Size-aware relative values for the original parallel queueing system with RR routing

State of the parallel queueing system

- Let **z** denote the current state of the queueing system,

$$\mathbf{z} = ((q_1, u_1), \ldots, (q_m, u_m))$$

- Here q_i refers to the queue currently in phase i and u_i to its backlog
- Let $v(\mathbf{z})$ denote the corresponding value function,

$$v(\mathbf{z}) = v_s(\mathbf{z}) + v_h(\mathbf{z})$$

Value function related to service costs

- Corollary 1:

$$v_s(\mathbf{z}) = \frac{1}{2m}\sum_{i=1}^{m}(2i-m-1)E[S^{(q_i)}]$$

Value function related to virtual waiting costs

- Corollary 2:

$$v_b(\mathbf{z}) = \sum_{i=1}^{m} v_i^{(q_i)}(u_i)$$

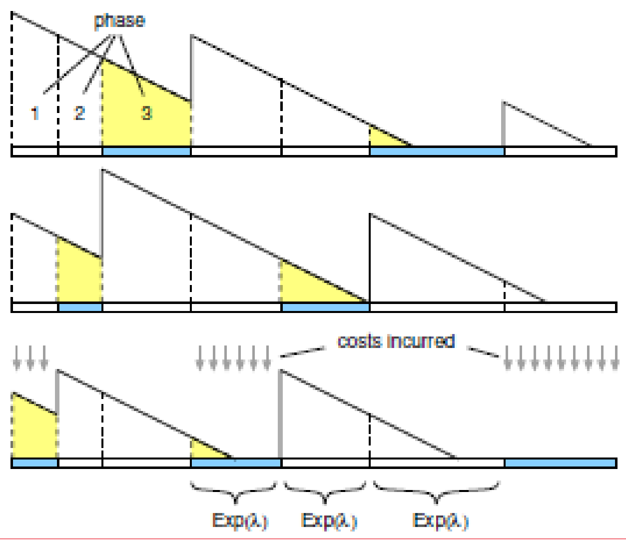

Relative values related to (real) waiting costs

- Assume:
  - $X_k \sim X$
  - $H_k = 1$
- Proposition 4:

$$v_w(\mathbf{z}) - v_w(\mathbf{0}) = \lambda(v_b(\mathbf{z}) - v_b(\mathbf{0}))$$

- Proof (idea): PASTA

Relative values related to holding costs

- Assume:
  - $X_k \sim X$
  - $H_k \sim H$
- Corollary 5:

$$v_h(\mathbf{z}) - v_h(\mathbf{0}) = \lambda(v_b(\mathbf{z}) - v_b(\mathbf{0}))E[H]$$

- Proof (idea): W_{jk} and H_{jk} independent for FCFS

Part IV
FPI-RR dispatching policy

FPI-RR dispatching policy

- Action a determines
  - the queue k to which the new job is assigned
  - the phases of all queues
- In state $\mathbf{z}$, FPI-RR assigns job j according to action a^*,

$$a^* = \arg\min_{a} \{ \underbrace{s_{j,k(a)} + w_{k(a)}(\mathbf{z}) h_{j,k(a)}}_{\text{immediate costs}} + \underbrace{v(\mathbf{z} \oplus (a, x_{j,k(a)})) - v(\mathbf{z})}_{\text{future add. costs}} \}$$

- Remark:
  - FPI-RR is size- and cost-aware utilizing the exact information on the arriving job (x, h, s) and the state of the system ($\mathbf{z}$)

Dispatching policies for M/D/2 with $S_1 = 1$, $S_2 = 4$, $H = 1$, $\rho = 0.4$

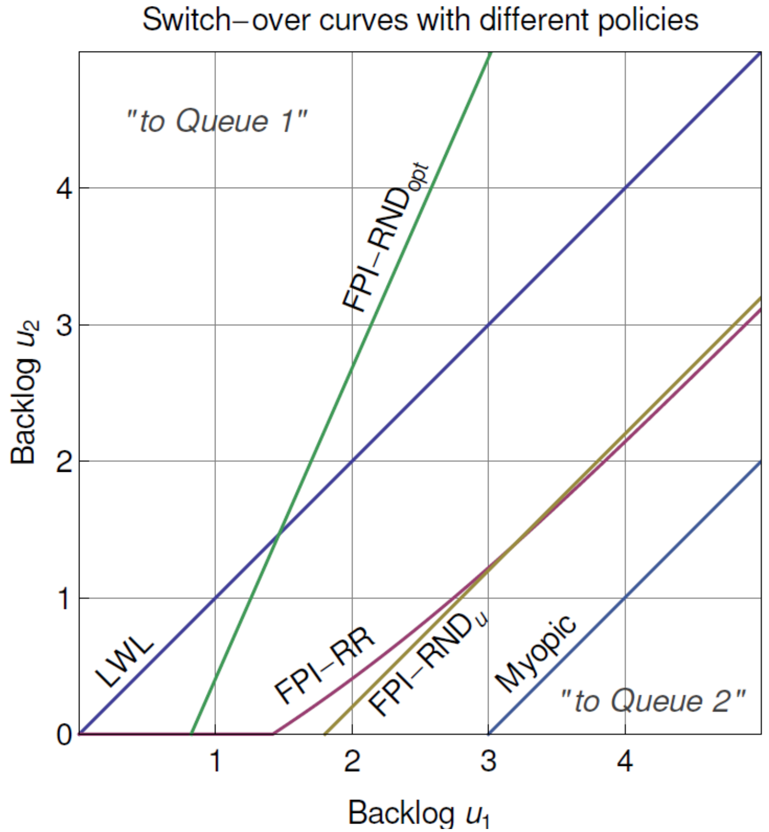

NOTE:

In this case:
LWL = RR

Mean holding costs for M/G/2 with S = 0 and H ~ Exp(1)

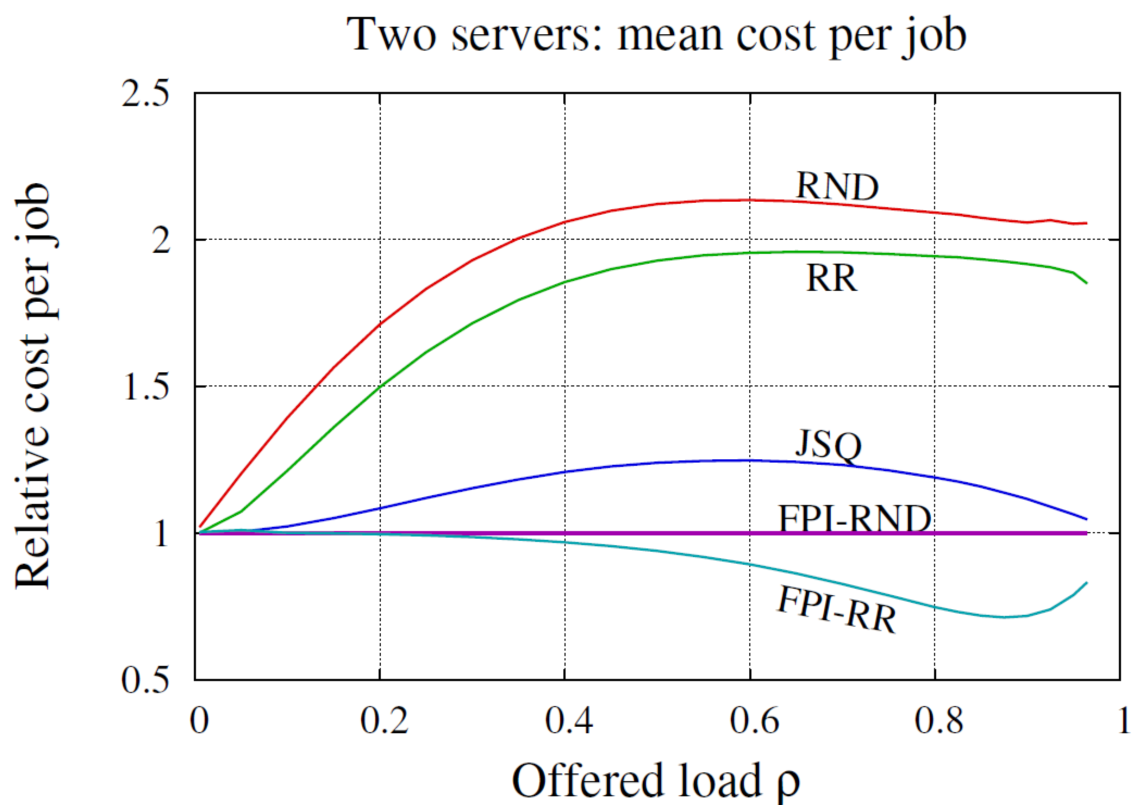

NOTE:

In this case: LWL = FPI-RND

Summary

- Dispatching problem in parallel queues with job- and queue-specific service and holding costs
- Value functions characterized for the Erl(m,λ)/G/1 queue
- Size-aware relative values determined for the original parallel queueing system with Poisson arrivals and RR routing
- FPI-RR dispatching policy derived to improve RR
- Results possible to be generalized to more complex (deterministic) routing patterns than RR

The End