Optimal Control of Polling Systems with Large Switchover Times

Yue Hu (Columbia Business School)

Joint work with
Jing Dong (Columbia Business School)
and
Ohad Perry (Northwestern University)

The Model

A polling system is a stochastic queueing network where several queues are served by one switching server

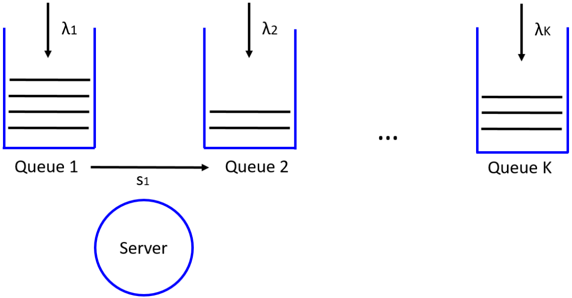

The Model

A polling system is a stochastic queueing network where several queues are served by one switching server

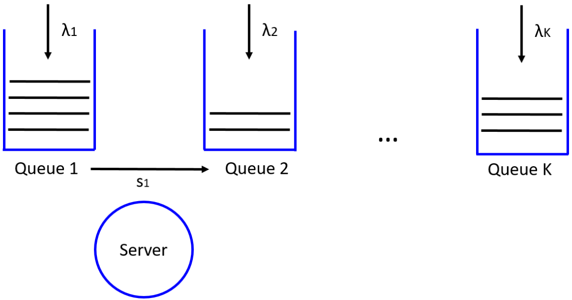

- A single server switching between $K \geq 2$ queues
- The server attends to the queues in a periodic cycle of a stages defined by the polling table $p : \{1, ..., a\} \rightarrow \{1, ..., K\}$ $(a \geq K)$
  - $p(j) = i$ – the server serves queue i at stage j within a cycle

The Model

A polling system is a stochastic queueing network where several queues are served by one switching server

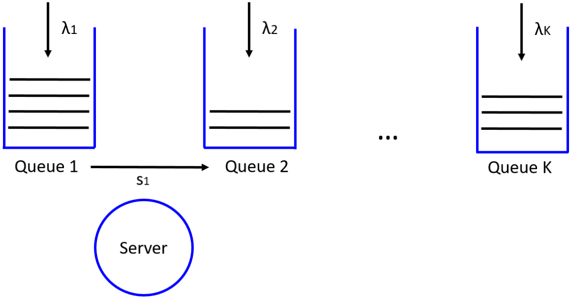

- Stationary arrival process of jobs with rate λ_i to queue i
- IID service times with mean $1/\mu_i$ at queue $i \in \{1, \ldots, K\}$
- Necessary stability condition: $\sum_{i=1}^{K} \lambda_i/\mu_i < 1$
- Large switchover time s_i is incurred when server switches from queue i

Stochastic Dynamics

- $Q_i(t)$ – number in queue i at time t, $1 \leq i \leq K$
- $Z(t)$ – the stage of the server at time t:
  $Z(t) = j$ – server is at stage j, $j = 1, ..., a$
  $Z(t) = \ominus_j$ – server is switching from stage j to stage $j + 1$ (modulo a)

The stochastic dynamics are characterized via the process

$$X(t) := (Q_i(t), Z(t); i = 1, \ldots, K), \quad t \geq 0$$

Remark: Given a routing order and service policy, $X(t)$ is a well-defined stochastic process

Applications

Polling systems are used to model computer, communication, and production systems

Stochastic economic lot scheduling: make-to-stock production of multiple products on a single machine

Related Literature

Polling System:

- Zero Switchover Time: Coffman et al. (1995), Remerova et al. (2014)
- Small Switchover Time: Boxma et al. (1990), Fricker and Jaibi (1993), Borst and Boxma (1997), Coffman et al. (1998), Olsen and Van der Mei (2003, 2005)
- Large Switchover Time: Van der Mei (1999), Olsen (2001), Lan and Olsen (2006)

Hybrid Dynamical Systems:

Matveev et al. (2016), Feoktistova et al. (2012), Perry and Whitt (2016), Dong and Perry (2018)

Goal

Our goal is to find a control that minimizes long run average holding costs among all “admissible” controls that are

- non-idling
- discrete-Markov (Given the queue length at the polling instant, the service policies do not depend on the history of the service process up to that instant)

Remark:

- $X(t)$ is not necessarily a markov process under an admissible control
- It is prohibitively hard to solve the optimal control problem exactly.
- We carry out optimality analysis in an asymptotic sense.

A Fluid Approximation to Simplify the Problem

Consider a hybrid dynamic system (HDS) characterized by

$$x(t) = (q(t), z(t)),$$

where $x(t)$ is a hybrid of the piece-wise linear continuous number-in-system process $q(t) \in [0, \infty)$ and the discrete server-position process $z(t) \in \{j, \ominus_j, 1 \leq j \leq a\}$, characterized via

$$\dot{q}_i(t) = (\lambda_i - \mu_i)\mathbf{1}_{\{p(z(t))=i, q(t)>0\}} + \lambda_i \mathbf{1}_{\{p(z(t)) \neq i\}}, \quad 1 \leq i \leq K$$

A Fluid Approximation to Simplify the Problem

Consider a hybrid dynamic system (HDS) characterized by

$$x(t) = (q(t), z(t)),$$

where $x(t)$ is a hybrid of the piece-wise linear continuous number-in-system process $q(t) \in [0, \infty)$ and the discrete server-position process $z(t) \in \{j, \ominus_j, 1 \leq j \leq a\}$, characterized via

$$\dot{q}_i(t) = (\lambda_i - \mu_i)\mathbf{1}_{\{p(z(t))=i, q(t)>0\}} + \lambda_i \mathbf{1}_{\{p(z(t))\neq i\}}, \quad 1 \leq i \leq K$$

Remark:

- Under a fixed routing order and service policy, $x(t)$ is well-defined

A Fluid Approximation to Simplify the Problem

Consider a hybrid dynamic system (HDS) characterized by

$$x(t) = (q(t), z(t)),$$

where $x(t)$ is a hybrid of the piece-wise linear continuous number-in-system process $q(t) \in [0, \infty)$ and the discrete server-position process $z(t) \in \{j, \ominus_j, 1 \leq j \leq a\}$, characterized via

$$\dot{q}_i(t) = (\lambda_i - \mu_i)\mathbf{1}_{\{p(z(t))=i, q(t)>0\}} + \lambda_i \mathbf{1}_{\{p(z(t)) \neq i\}}, \quad 1 \leq i \leq K$$

Fluid Optimal Control Problem: To find a control that minimizes

$$c := \limsup_{t \to \infty} \frac{1}{t} \int_0^t \sum_{i=1}^{K} \psi_i(q_i(s))ds,$$

where $\psi_i(\cdot)$ is the cost rate for queue i

The Road Map

A Road Map of Our Approach:

- Find the optimal periodic equilibrium among all the possible periodic equilibria of the HDS
- Design a fluid control that achieves the optimal equilibrium
- “Interpret” the fluid control into a control for the stochastic system
- Prove: The control is asymptotically optimal

Characterization of Periodic Equilibria

E.x. periodic equilibrium with two queues and server routing $\{1, 2|1, 2| \ldots\}$

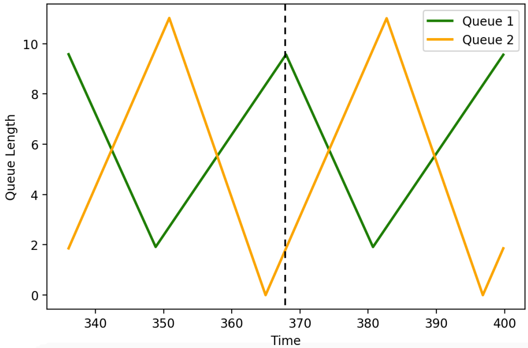

Characterization of Periodic Equilibria

E.x. periodic equilibrium with two queues and server routing $\{1, 2|1, 2|\ldots\}$

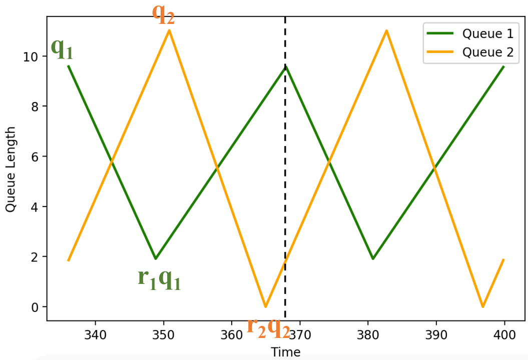

Fixed-Proportion Reduction Control (FPR)

E.x. periodic equilibrium with two queues and server routing $\{1,2|1,2|\ldots\}$

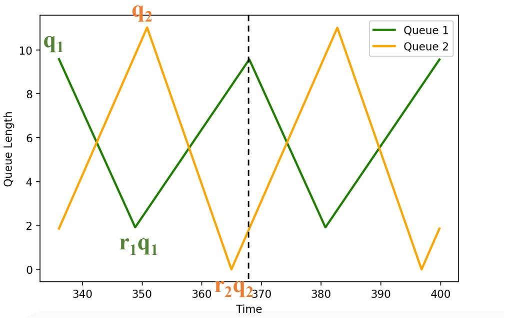

The FPR Control (Definition):

- For fixed $\{r_j \in [0,1],\ j = 1, ..., a\}$, if $p(j) = i$, then reduce q_i (the value of queue i at the beginning of service) to $r_i q_i$
- After serving station j, switch to station $j+1$ in the polling table

Fixed-Proportion Reduction Control (FPR)

E.x. periodic equilibrium with two queues and server routing $\{1, 2|1, 2|\dots\}$

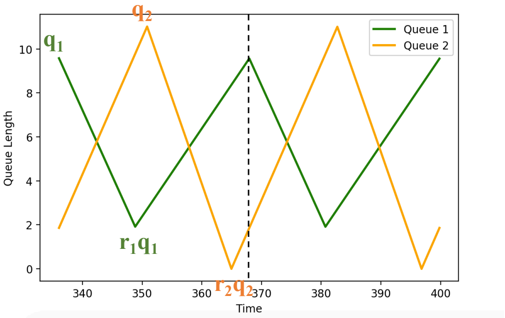

<u>Remark:</u> Any periodic equilibrium can be translated to some control parameters of the FPR

Fixed-Proportion Reduction Control (FPR)

E.x. periodic equilibrium with two queues and server routing $\{1, 2|1, 2|\dots\}$

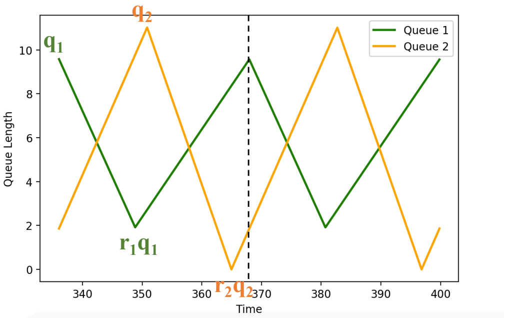

Theorem

Independently of initialization,

- *the HDS converges to a unique periodic equilibrium under FPR with any non-trivial control parameters*
- *the HDS converges to the desired equilibrium under the translated FPR*

The Optimal Control

There is a bijective mapping from the possible periodic equilibria of the HDS to the FPR control.

Therefore, minimizing the long-run average cost c is equivalent to finding the optimal FPR control:

$$\min_{\{r_j, j=1,\ldots,a\}} \quad \frac{1}{\tau}\int_0^{\tau}\sum_{i=1}^{K}\psi_i\left(q_i(s,r)\right)ds \quad s.t. \quad r_j \in [0,1], \quad j=1,\ldots,a,$$

where τ is the equilibrium cycle length

The Optimal Control

There is a bijective mapping from the possible periodic equilibria of the HDS to the FPR control.

Therefore, minimizing the long-run average cost c is equivalent to finding the optimal FPR control:

$$\min_{\{r_j, j=1,\dots,a\}} \quad \frac{1}{\tau} \int_0^{\tau} \sum_{i=1}^{K} \psi_i \left(q_i(s, r) \right) ds \quad s.t. \quad r_j \in [0, 1], \quad j = 1, \dots, a,$$

where τ is the equilibrium cycle length

- E.x.1: Under routing $\{1, 2, 3|1, 2, 3| \dots\}$, optimal ratio is $\{0\%, 0\%, 0\%\}$
- E.x.2: Under routing $\{1, 3, 2, 3, 2|1, 3, 2, 3, 2| \dots\}$, optimal ratio is $\{0\%, 48.86\%, 0\%, 0\%, 0\%\}$

$\lambda_i = 2, \mu_i = 8, s_i = 2, i = 1, 2, 3$
linear cost rates $\psi_1(\cdot) = 1\cdot, \psi_2(\cdot) = 4\cdot, \psi_3(\cdot) = 1\cdot$

The Optimal Control

There is a bijective mapping from the possible periodic equilibria of the HDS to the FPR control.

Therefore, minimizing the long-run average cost c is equivalent to finding the optimal FPR control:

$$\min_{\{r_j, j=1,\dots,a\}} \quad \frac{1}{\tau}\int_0^{\tau} \sum_{i=1}^{K} \psi_i \left(q_i(s,r)\right) ds \quad s.t. \quad r_j \in [0,1], \quad j = 1, \dots, a,$$

where τ is the equilibrium cycle length

Proposition

It is optimal to exhaust each queue at least once during a cycle

Corollary

In the case of cyclic routing, the exhaustive policy is optimal

The Road Map

The Road Map Revisited:

- Find the optimal periodic equilibrium among all the possible periodic equilibria of the HDS
- Design a fluid control that achieves the optimal equilibrium
- "Interpret" the fluid control into a control for the stochastic system
  - Reduce queue i from q_i to $\lceil r_j q_i \rceil$ if the server is at queue i in stage j
- Prove: The control is asymptotically optimal

Asymptotic Optimality

Consider a sequence of systems indexed by n under an admissible control π^n

- Large switchover times: switchover times satisfy $s_i^n = ns_i$ in n^{th} system
- Temporal and spacial scalings: $\bar{X}^n_{\pi^n}(t) := (Q_{\pi^n}(nt)/n, Z_{\pi^n}(nt))$
- Scaled long-run average costs: $\bar{C}^n_{\pi^n} = \lim_{t\to\infty} \frac{1}{t}\int_0^t \psi(\bar{Q}^n_{\pi^n}(s))ds$

Asymptotic Optimality

Consider a sequence of systems indexed by n under an admissible control π^n

- Large switchover times: switchover times satisfy $s_i^n = ns_i$ in n^{th} system
- Temporal and spacial scalings: $\bar{X}^n_{\pi^n}(t) := (Q_{\pi^n}(nt)/n, Z_{\pi^n}(nt))$
- Scaled long-run average costs: $\bar{C}^n_{\pi^n} = \lim_{t\to\infty} \frac{1}{t} \int_0^t \psi(\bar{Q}^n_{\pi^n}(s))ds$

Let π_*^n denote the optimal fluid FPR control translated for the n^{th} stochastic system, and c_* denote the optimal fluid objective value

Asymptotic Optimality

Consider a sequence of systems indexed by n under an admissible control π^n

- Large switchover times: switchover times satisfy $s_i^n = ns_i$ in n^{th} system
- Temporal and spacial scalings: $\bar{X}^n_{\pi^n}(t) := (Q_{\pi^n}(nt)/n, Z_{\pi^n}(nt))$
- Scaled long-run average costs: $\bar{C}^n_{\pi^n} = \lim_{t\to\infty} \frac{1}{t}\int_0^t \psi(\bar{Q}^n_{\pi^n}(s))ds$

Let π^n_* denote the optimal fluid FPR control translated for the n^{th} stochastic system, and c_* denote the optimal fluid objective value

Theorem (Asymptotic Optimality)

- *Under any sequence of policies π^n, $\liminf_{n\to\infty} \bar{C}^n_{\pi^n} \geq c_*$*
- *Under π^n_*, $\bar{C}^n_{\pi^n_*} \Rightarrow c_*$ as $n \to \infty$ (The FPR is asymptotically optimal)*

Asymptotic Optimality

Consider a sequence of systems indexed by n under an admissible control π^n

- Large switchover times: switchover times satisfy $s_i^n = ns_i$ in n^{th} system
- Temporal and spacial scalings: $\bar{X}^n_{\pi^n}(t) := (Q_{\pi^n}(nt)/n, Z_{\pi^n}(nt))$
- Scaled long-run average costs: $\bar{C}^n_{\pi^n} = \lim_{t\to\infty} \frac{1}{t} \int_0^t \psi(\bar{Q}^n_{\pi^n}(s))ds$

Let π^n_* denote the optimal fluid FPR control translated for the n^{th} stochastic system, and c_* denote the optimal fluid objective value

Theorem (Asymptotic Optimality)

- *Under any sequence of policies* π^n, $\liminf_{n\to\infty} \bar{C}^n_{\pi^n} \geq c_*$
- *Under* π^n_*, $\bar{C}^n_{\pi^n_*} \Rightarrow c_*$ *as* $n \to \infty$ *(The FPR is asymptotically optimal)*

$\Rightarrow$ For cyclic routing, the exhaustive policy is asymptotically optimal

Asymptotic Optimality (Proof)

$$\begin{array}{ccc}
\bar{Q}^n(t) & \xrightarrow{t \to \infty} & \bar{Q}^n_*(t) \\
{\scriptstyle n \to \infty}\downarrow & & \downarrow{\scriptstyle n \to \infty} \\
q(t) & \xrightarrow[t \to \infty]{} & q_*(t)
\end{array}$$

Theorem (Limit Interchange)

Under any FPR control parameters,

$$\bar{C}^n = \lim_{t\to\infty} \frac{1}{t} \int_0^t \psi\left(\bar{Q}^n(s)\right) ds \Rightarrow c \text{ as } n \to \infty$$

In particular, $\bar{C}^n_{\pi^n_*} \Rightarrow c_*$ *as* $n \to \infty$ *under* π^n_*

Summary

- We considered the optimal-control problem of a polling system with large switchover times when long-run average holding costs are to be minimized
- We proposed the Fixed-Proportion Reduction Control (FPR):
  - all possible periodic equilibria are achievable under FPR
  - stable under all control parameters
- We find the optimal FPR and prove that the control is asymptotically optimal under large-switchover-time scaling
- In the special case of cyclic routing, the exhaustive policy is optimal

Thank You