IBM Research

Optimal Control in Fluid Models of $n \times n$ Input-Queued Switches under Linear Fluid-Flow Costs

Y. Lu, M.S. Squillante, T. Suk
Mathematical Sciences
IBM Thomas J. Watson Research Center

Input-Queued Switch

- Input-queued switches are widely used in computer and communication networks

- 3×3 Input-queued switch:

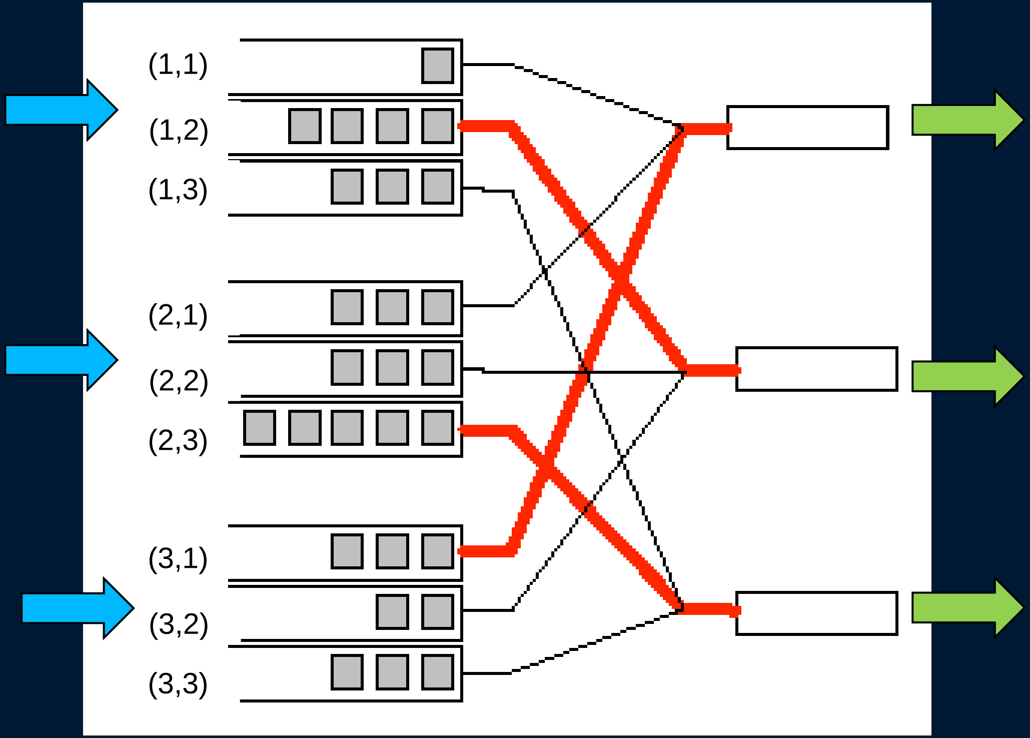

Input-Queued Switch

- Consider input-queued switch with n input ports and n output ports
- Each input port has queue associated with every output port that stores packets waiting to be transmitted
- Simultaneous transmission of packets is possible only from certain subsets of the queues, as defined by following constraints:
  - Every input port can transmit at most one packet
  - Every output port can receive at most one packet
- We call the subsets of queues that satisfies these constraints basic schedules

Analysis of the Input-Queued Switches

- Main focus of previous research: Throughput optimality
  - E.g.: Tassiulas and Ephremides (1992); McKeown, Anantharam, and Walrand (1996)
- Study of delay optimality focused on MaxWeight and heavy traffic regime
  - E.g.: Kang and Williams (2012); Maguluri and Srikant (2016); Lu, Maguluri, Squillante, and Suk (2018b)
- Optimal policy obtained in 2x2 case, reveals some different structures (e.g. switching curve)
  - Lu, Maguluri, Squillante, and Suk (2018a) for original stochastic system under general linear-cost objective function
  - These optimal results and structures can be generalized to the $n \times n$ switch only in special cases, and not in general
- Fluid models
  - E.g.: Shah and Wischik (2012) and, more recently, Sharifnassab, Tsitsiklis, and Golestani (2020) on fluid models under MaxWeight
  - General linear fluid flow cost structures

Overview

- Stochastic Model of Input-Queued Switch
- Fluid Model for Input-Queued Switch and Optimal Control Problem
  - Difficulty of Optimal Control Problem
- Optimal Control Algorithm
  - Critical Threshold
- Main Theoretical Results
  - Stability
  - Optimality
- Computational Experiments
- Conclusion

Stochastic Model

- Queue with input port i, output port j is indexed by $(i, j) \in \mathbb{J} := [n] \times [n]$
- Time is slotted and denoted by a nonnegative integer $t \in \mathbb{Z}_+ := \{0, 1, \ldots\}$
- Service time of packets is 1 time unit
- At each time t, scheduling policy selects a basic schedule such that packet from nonempty queue in the schedule is served
- Basic schedule formally depicted by n^2-dimensional binary vector $s = (s_{ij})_{ij \in [n]}$ such that $s_{ij} = 1$ if queue (i, j) is in schedule, and $s_{ij} = 0$ otherwise
- Set of all basic schedules $\mathbb{I}$:

$$\mathbb{I} = \left\{ s \in \{0,1\}^{\mathbb{J}} : \sum_{i \in [n]} s(i,j) \leq 1, \sum_{j \in [n]} s(i,j) \leq 1, \forall i, j \in [n] \right\}.$$

Dynamics of Stochastic Model

- $\mathcal{Q}_{ij}(t)$: length of queue (i,j) at beginning of t-th slot; $\mathcal{Q}(t) = \{\mathcal{Q}_{ij}(t)\}$
- $\mathcal{A}_{ij}(t) \in \mathbb{Z}_+$: number of arrivals to queue (i,j) up to time t, where
  - $\{\mathcal{A}_{ij}(t+1) - \mathcal{A}_{ij}(t)\}$ i.i.d. with $\mathbb{E}[\mathcal{A}(t+1) - \mathcal{A}(t)] = \lambda$,
  - arrival rate vector $\boldsymbol{\lambda} \in \mathbb{R}_+^{|\mathbb{J}|}$
- $\mathcal{D}_s(t)$: Cumulative number of time slots devoted to basic schedule s until t:

$$\|\mathcal{D}(t)\| = t, \qquad \|\mathcal{D}(t+1) - \mathcal{D}(t)\| = 1$$

- Queueing dynamics:

$$\mathcal{Q}(t) = \mathcal{Q}_0 + \mathcal{A}(t) - \mathcal{D}(t)A$$

where A is the $|\mathbb{I}| \times |\mathbb{J}|$–dimensional binary schedule-queue adjacency matrix:

$$A_{s,(i,j)} = s_{ij}$$

Input-Queued Switch Scheduling: Fluid Model

Consider r-scaled process: $(\mathcal{Q}^r(t), \mathcal{A}^r(t), \mathcal{D}^r(t)) := \left(\frac{1}{r}\mathcal{Q}(rt), \frac{1}{r}\mathcal{A}(rt), \frac{1}{r}\mathcal{D}(rt)\right)$

- $\lim_{r\to\infty} \sup_{0\le t\le T} \|\mathcal{A}^r(t) - \lambda t\| = 0, \quad \mathcal{D}_{ij}^r(t') - D_{ij}^r(t) \le (t - t')$

Convergent subsequence of $\mathcal{Q}^r(t)$ converges to Fluid Model $q(t)$ such that

$$\dot{q}(t) = \lambda - \sigma(t)A,$$

$$q(t) \ge 0, \qquad \|\sigma(t)\| = 1, \qquad \sigma(t) \ge 0$$

Fluid-level schedule is a convex combination of basic schedules

$\big(q(t), \sigma(t)\big)$: Fluid-level admissible pair

$\sigma(t)$: Fluid-level admissible policy

Fluid Model Optimal Control Problem

- $c = \{c_{ij}\}$: cost coefficient vector
- Define total discounted queue-length cost over the entire time horizon under a fluid-level admissible policy $\{\sigma(t) : t \in \mathbb{R}_+\}$ with initial state $\boldsymbol{q}_0$:

Fluid Optimal Control Problem

$$minimize \int_0^\infty e^{-\beta t} c \cdot q(t) dt$$

$$\dot{q}(t) = \lambda - \sigma(t) A$$

$$q(t) \geq 0$$

$$\sigma(t) \geq 0$$

$$\|\sigma(t)\| = 1$$

Difficulty of Optimal Control Problem

- While optimal control framework enables with relative ease derivation of optimal policies for fluid models of basic queueing networks, situation for input-queued switches is quite different and much more difficult
- For example, arrival rate vector $\boldsymbol{\lambda}$ and initial queue length $\boldsymbol{q}_0$ s.t. $\lambda_{ij} = 0, \forall i \in [n], \forall j \in [n] \setminus \{1\}$, then equivalent to n parallel queues with one server
In this case, $c\mu$-policy well-known to be optimal policy that minimizes discounted total cost over infinite horizon in both stochastic and fluid model
However, $c\mu$-policy is not always stable even in the fluid limit model
- As another example, MaxWeight Scheduling Algorithm is stable

Example: Unstable Case of $c\mu$-policy

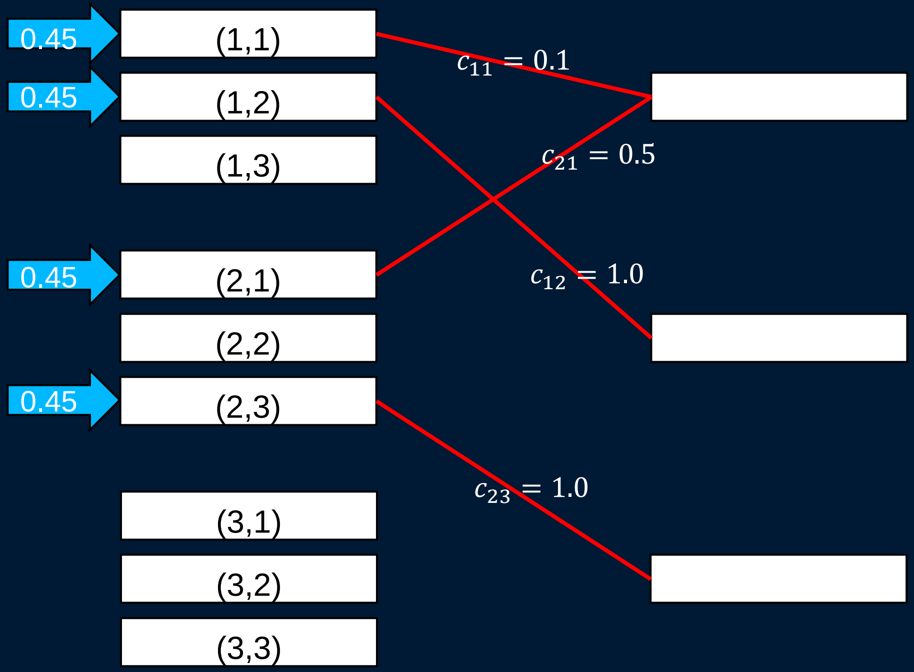

- Maximum Basic Schedules: {(1,1),(2,3)}, {(1,2),(2,1)}, {(1,2),(2,3)}

- $c\mu$-policy:

$$0.45 \times \{(1,2),(2,3)\}$$
$$+0.45 \times \{(2,1)\}$$
$$+0.1 \times \{(1,1)\}$$

- $c\mu$-policy is not always stable even in fluid limit model

Optimal Control Algorithm: Critical Threshold

- For q and $\tau \in \mathbb{R}_+$, define

Associated LP(q, τ)

$$\text{maximize } (Ac) \cdot \sigma - \tau \|\sigma\|$$
$$(\sigma A)_{ij} \le \lambda_{ij} \ \forall (i,j) \text{ with } q_{ij} = 0$$
$$\sigma \ge 0$$

Fluid Optimal Control Problem

$$\text{minimize} \int_0^\infty e^{-\beta t} c \cdot q(t) dt$$
$$\dot{q}(t) = \lambda - \sigma(t) A$$
$$q(t) \ge 0$$
$$\sigma(t) \ge 0$$
$$\|\sigma(t)\| = 1$$

- $LP(q, \tau)$ maximizes the weighted outflow, subjective to feasibility constraint
- τ : the multiplier of the constraint $\|\sigma(t)\| = 1$

Optimal Control Algorithm: Critical Threshold

- For q and $\tau \in \mathbb{R}_+$, define

Associated LP(q, τ)

$$maximize\ (Ac) \cdot \sigma - \tau \|\sigma\|$$
$$(\sigma A)_{ij} \leq \lambda_{ij}\ \forall (i,j)\ with\ q_{ij} = 0$$
$$\sigma \geq 0$$

Fluid Optimal Control Problem

$$minimize \int_0^\infty e^{-\beta t} c \cdot q(t) dt$$
$$\dot{q}(t) = \lambda - \sigma(t) A$$
$$q(t) \geq 0$$
$$\sigma(t) \geq 0$$
$$\|\sigma(t)\| = 1$$

- If $\exists \sigma^*$ an optimal solution such that $\|\sigma^*\| = 1$, τ is called a critical threshold of state q

Theorem

There always exists a critical threshold for any state q.

The critical thresholds can be found via a set of search algorithms.

Main Results: Optimal Control Algorithm

Algorithm 4 Optimal Control Algorithm for initial state $\boldsymbol{q}_{t=0}$

1. Set $k = 0$, $t_0 = 0$, and $\boldsymbol{q}_0^* = \boldsymbol{q}_{t=0}$
2. **while** $t_k < \infty$ **do**
3. Let τ_k be the output of Algorithm 3 with input $\boldsymbol{q} = \boldsymbol{q}_{t_k}^*$
4. Let γ_k be the optimal value of Problem $(P_{\boldsymbol{q},\tau})$ with $\boldsymbol{q} = \boldsymbol{q}_{t_k}^*$ and $\tau = \tau_k$
5. Find a point $\boldsymbol{\nu}_k \in \mathbb{Q}(\boldsymbol{q}_{t_k}^*, \tau_k, \gamma_k)$ in (8)
6. Define $\boldsymbol{\mu}^* \in \mathbb{R}^{\mathbb{I}}$ by
$$\mu^*(s) = \begin{cases} \nu_k(s) & \text{if } s \in \mathbb{I}_{\tau_k} \\ 0 & \text{otherwise} \end{cases}$$
7. Set
$$t_{k+1} = t_k + \min\left\{ \frac{q_{t_k}(\rho)}{(\mu^* A)(\rho) - \lambda(\rho)} : \rho \in \mathbb{J} \setminus \mathbb{J}_{\boldsymbol{q}_{t_k}^*},\ (\mu^* A)(\rho) - \lambda(\rho) > 0 \right\}$$
8. Set $\boldsymbol{\mu}^*(t) = \boldsymbol{\mu}^*$ for $t \in [t_k, t_{k+1})$ and $\boldsymbol{q}_t^* = \boldsymbol{q}_{t_k}^* + (t - t_k)\boldsymbol{\lambda} - (t - t_k)\boldsymbol{\mu}^* \boldsymbol{A}$ for $t \in [t_k, t_{k+1}]$
9. Set $k = k + 1$

Algorithm 3 Algorithm to find a critical threshold at state $\boldsymbol{q}$

Input: State $\boldsymbol{q}$ **Output:** a critical threshold $\tau = \tau(\boldsymbol{q})$

1. Set m be the output of Algorithm 1 with input $\boldsymbol{q}$
2. **if** $m > 0$ **then**
3. **return** τ_m
4. **else**
5. **return** the output of Algorithm 2 with input $l = -m$

- Algorithm 4: optimal control algorithm
  - Starting at any q, find the critical threshold τ
  - Follow the allocation rule from LP($\boldsymbol{q}$,$\boldsymbol{\tau}$) until one of the queues reaches zero;
  - Repeat
- Algorithm 3 is the mega-algorithm for using Algorithm 1 and 2 to obtain the critical threshold

Critical Threshold: Example

- If $q_{ij} \neq 0$ for all i, j

Associated LP$(\boldsymbol{q}, \boldsymbol{\tau})$

$$maximize\ (Ac) \cdot \sigma - \tau \|\sigma\|$$
$$\sigma \geq 0$$

- Critical threshold $\tau = \max\{c \cdot s : s \in \mathbb{I}\}$ and $\sigma^* = \operatorname{argmax}\{c \cdot s : s \in \mathbb{I}\}$
- Coincides with the $c\mu$-rule
- In the 3x3 case where $q_{12} = q_{21} = 0$ and $q_{11} = q_{23} > 0$, the critical threshold is given by $\tau = 0$ and the optimal policy is given by: 0.45 for (1,2) and (2,1); 0.55 for (1,1) and (2,3)

Main Results: Optimal Control Algorithm

- At state q, use optimal solution σ^* of associated LP such that $||\sigma^*|| = 1$ for critical threshold τ

Theorem (Stability or Throughput-Optimality)

If $\sum_i \lambda_{ij} < 1$ and $\sum_j \lambda_{ij} < 1$ for all i, j, the above set of algorithms empty the system in finite time

Theorem (Optimality)

If $\sum_i \lambda_{ij} < 1$ and $\sum_j \lambda_{ij} < 1$ for all i, j, the above set of algorithms provides an optimal solution to the Fluid Optimal Control Problem

Main Results: Optimal Control Algorithm

- At state q, use optimal solution σ^* of associated LP such that $||\sigma^*|| = 1$ for critical threshold τ

Theorem (Stability or Throughput-Optimality)

If $\sum_i \lambda_{ij} < 1$ and $\sum_j \lambda_{ij} < 1$ for all i, j, the above set of algorithms empty the system in finite time

Main idea: Caratheodory's Theorem key to construct a Lyapunov function

Main Results: Optimal Control Algorithm

- At state q, use optimal solution σ^* of associated LP such that $||\sigma^*|| = 1$ for critical threshold τ

Theorem (Optimality)

If $\sum_i \lambda_{ij} < 1$ and $\sum_j \lambda_{ij} < 1$ for all i, j, the above set of algorithms provides an optimal solution to the Fluid Optimal Control Problem

Main idea: verify the necessary and sufficient condition for Pontryagin's Maximum Principle

Necessary and Sufficient Conditions

Admissible policy $\sigma^*(t)$ is optimal solution to Fluid Optimal Control Problem

if $\exists p(t), \eta(t)$ such that

- $\sigma^*(t) \in \operatorname{argmax}\{\sigma A p(t) : \sigma \geq 0, \|\sigma\| = 1\}$
- $\dot{p}(t) - \beta\, p(t) = c - \eta(t)$
- $q^*(t) \cdot \eta(t) = 0, q^*(t) \geq 0, \eta(t) \geq 0$
- $\liminf\, p(t) \cdot \left(q^*(t) - q(t)\right) \geq 0$ for any fluid model $q(t)$

$\eta(t)$: solution to the dual problem of associated LP

$$p(t) \coloneqq \int_t^T e^{(T-t')}\left(c - \eta(t')\right) dt$$

Computational Experiments

- Compare through simulations performance of our optimal control algorithm with that of $c\mu$-rule and max-weight scheduling algorithm in fluid model

- Fix number of input and output ports to be $n \in \mathbb{Z}_+$ and fix throughput $\kappa \in (0,1)$

- For $1 \leq i, j \leq n$, randomly generate costs $c_{ij} \in (0,1)$ and arrival rates $\lambda_{ij} \in (0,1)$ such that

$$\max \left\{ \sum_{k=1}^{n} \lambda(i,k), \sum_{k=1}^{n} \lambda(k,j) \ : \ i,j \in [n] \right\} = \kappa. \qquad (24)$$

We also choose an initial queue length to be an integer between 1 and 100 uniformly at random for each $(i,j) \in [n] \times [n]$.

Computational Experiments

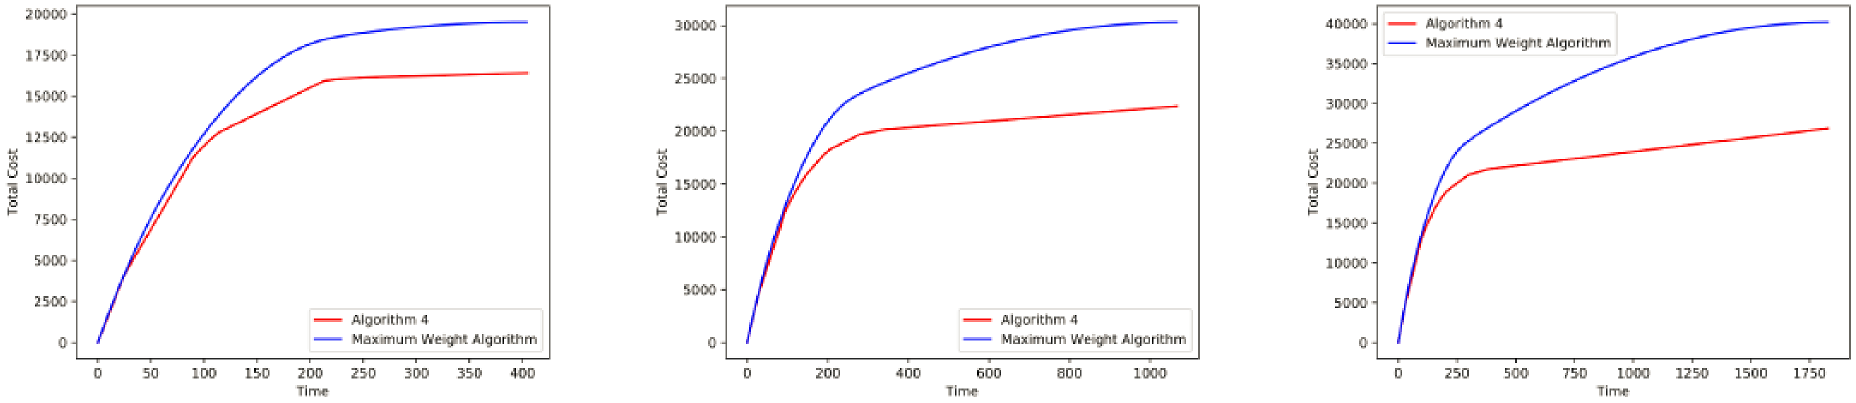

(a) $\kappa = 0.70$, **Relative gap is** 19% **(b)** $\kappa = 0.90$, **Relative gap is** 35% **(c)** $\kappa = 0.95$, **Relative gap is** 50%

Figure 1: Total costs of Algorithm 4 and Max-Weight Algorithm

Computational Experiments

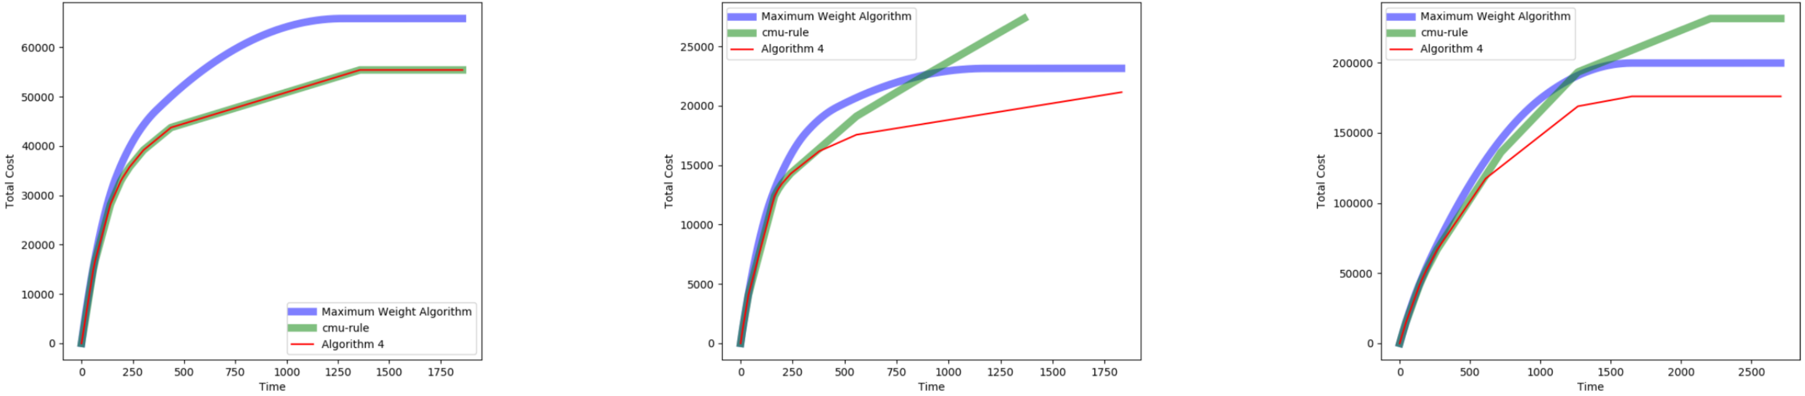

(a) Algorithm 4 coincides $c\mu$-rule **(b) Unstable $c\mu$-rule** **(c) Stable but not optimal $c\mu$-rule**

Figure 3: Performance Comparisons of Total Costs under Optimal Policy (Algorithm 4) and $c\mu$-rule

Conclusion

- Considered fluid model of general $n \times n$ input-queued switches where each fluid flow has associated cost
- Derived optimal scheduling control policy under general linear objective function based on minimizing discounted fluid cost over infinite horizon
- Optimal policy coincides with $c\mu$-rule in certain parameter domains
- In general, optimal policy determined algorithmically by constrained flow maximization problem whose Lagrangian multipliers of some key network constraints were identified by set of carefully designed algorithms
- Computational experiments within fluid models of input-queued switches demonstrated significant benefits of our optimal scheduling policy over alternative policies such as the $c\mu$ and max-weight scheduling policies

IBM Research

IBM

Optimal Control of Fluid Models of Switched Networks

Y. Lu, M.S. Squillante, T. Suk
Mathematical Sciences
IBM Thomas J. Watson Research Center