Optimal Delay-Cost Scheduling Control in Fluid Models of General n×n Input-Queued Switches

Yingdong Lu, Mark S. Squillante, Tonghoon Suk
Mathematical Sciences, IBM Research
Yorktown Heights, NY 10598, USA

1. INTRODUCTION

Input-queued switch (IQS) architectures are widely used in modern computer/communication networks. The optimal scheduling control of these high-speed, low-latency networks is critical for our understanding of fundamental design and performance issues related to internet routers, cloud computing data centers, and high-performance computing. A large and rich literature exists around optimal scheduling in these systems. This includes the extensive study of IQSs as an important mathematical model for a general class of optimal scheduling control problems of broad interest.

Most of the previous research related to scheduling in IQSs has focused on throughput optimality, specifically establishing throughput optimality of the max-weight scheduling policy though with some recent research showing that max-weight scheduling is asymptotically optimal in heavy traffic for certain unit-cost queue length (QL) based objectives. However, the question of delay-optimal scheduling in IQSs remains open in general, as does the question of delay-optimal scheduling under more general objective functions.

To gain fundamental insights into these very difficult problems, we consider a fluid model of general $n \times n$ IQSs where each fluid flow has an associated cost and we derive an optimal scheduling control policy to an infinite horizon discounted control problem with a general linear objective function of fluid cost. Our optimal policy coincides with the $c\mu$-rule in certain parameter domains. More generally, due to the highly constrained structure of IQS networks, the optimal policy takes the form of the solution to a flow maximization problem, after we identify the Lagrangian multipliers of some key constraints through carefully designed algorithms. These theoretical results reflect the high complexity nature of IQSs, and are expected to be of interest more broadly.

We next describe our mathematical models and the optimal control problem. Then we present our analysis and results for optimal scheduling and related theoretical properties. Additional details and proofs are provided in [3].

2. MATHEMATICAL MODELS

2.1 Stochastic Models

The IQS of interest consists of n input ports and n output ports. For each pair $(i, j) \in \mathbb{J} := [n] \times [n]$, packets that needs to be transmitted from the i-th input port to the j-th output port are stored in a queue indexed by (i, j). Time is slotted by nonnegative integers and the length of queue $\boldsymbol{\rho} \in \mathbb{J}$ at the beginning of the t-th time slot is denoted by $\mathcal{Q}_t(\boldsymbol{\rho})$. External packets arrive at each queue according to an exogenous stochastic process. Let $\mathcal{A}_t(\boldsymbol{\rho}) \in \mathbb{Z}_+$ represent the number of arrivals to queue $\boldsymbol{\rho} \in \mathbb{J}$ up until time t. Assume that $\{\mathcal{A}_{t+1}(\boldsymbol{\rho}) - \mathcal{A}_t(\boldsymbol{\rho}) : t \in \mathbb{Z}_+, \boldsymbol{\rho} \in \mathbb{I}\}$ are independent random variables and that, for fixed $\boldsymbol{\rho} \in \mathbb{J}$, $\{\mathcal{A}_{t+1}(\boldsymbol{\rho}) - \mathcal{A}_t(\boldsymbol{\rho}) : t \in \mathbb{Z}_+\}$ are identically distributed with $\mathbb{E}[\mathcal{A}_{t+1}(\boldsymbol{\rho}) - \mathcal{A}_t(\boldsymbol{\rho})] =: \lambda(\boldsymbol{\rho})$.

During each time slot, packets in the queues can be simultaneously transmitted subject to: (1) At most one packet can be transmitted from an input port; (2) At most one packet can be transmitted to an output port. We denote the departure of packets from the queues during a time slot by an n^2-dimensional binary vector $\boldsymbol{s} = [s(i, j) : (i, j) \in \mathbb{J}]$ s.t. $s(i, j) = 1$ if a packet in queue (i, j) departs from the queue, and $s(i, j) = 0$ otherwise. We refer to $\boldsymbol{s}$ as a *basic schedule*, and let $\mathbb{I}$ denote the set of all basic schedules: $\mathbb{I} = \{\boldsymbol{s} \in \{0, 1\}^{\mathbb{J}} : \sum_{i\in[n]} s(i, j) \le 1, \sum_{j\in[n]} s(i, j) \le 1, \forall i, j \in [n]\}$.

Note that that the empty basic schedule $\boldsymbol{s}$, s.t. $s(i, j) = 0 \forall (i, j) \in \mathbb{J}$, is indeed a member of $\mathbb{I}$. For $\boldsymbol{s} \in \mathbb{I}$, let $\mathcal{D}_t(\boldsymbol{s})$ denote the cumulative number of time slots devoted to basic schedule $\boldsymbol{s}$ in the time interval $[0, t]$. We therefore have $\|\boldsymbol{\mathcal{D}}_t\|_1 = \sum_{\boldsymbol{s}\in\mathbb{I}} \mathcal{D}_t(\boldsymbol{s}) = t$ and $\|\boldsymbol{\mathcal{D}}_{t+1}\|_1 - \|\boldsymbol{\mathcal{D}}_t\|_1 = 1$ for every $t \in \mathbb{Z}_+$. From the description of arrivals and departures, we can see that $\boldsymbol{\mathcal{Q}}_t$ evolves according to the following dynamics $\boldsymbol{\mathcal{Q}}_t = \boldsymbol{\mathcal{Q}}_0 + \boldsymbol{\mathcal{A}}_t - \boldsymbol{\mathcal{D}}_t \boldsymbol{A}$, where $\boldsymbol{\mathcal{Q}}_0 = [\mathcal{Q}_0(\boldsymbol{\rho}) : \boldsymbol{\rho} \in \mathbb{J}]$ is the initial QLs and $\boldsymbol{A} \in \{0, 1\}^{\mathbb{I}\times\mathbb{J}}$ is the schedule-queue adjacency matrix s.t. $A(\boldsymbol{s}, \boldsymbol{\rho}) = s(\boldsymbol{\rho})$ for $\boldsymbol{s} \in \mathbb{I}$ and $\boldsymbol{\rho} \in \mathbb{J}$. We refer to a stochastic process $\{(\boldsymbol{\mathcal{Q}}_t, \boldsymbol{\mathcal{A}}_t, \boldsymbol{\mathcal{D}}_t) \in \mathbb{Z}_+^{\mathbb{J}} \times \mathbb{Z}_+^{\mathbb{J}} \times \mathbb{Z}_+^{\mathbb{I}} : t \in \mathbb{Z}_+\}$ that satisfies the above dynamics as a *discrete-time stochastic model for IQSs* with the (random) initial state $\boldsymbol{\mathcal{Q}}_0 \in \mathbb{Z}_+^{\mathbb{J}}$.

2.2 Fluid Models

We introduce a continuous-time deterministic process representing our fluid models for IQSs through Definition 2.1. The basic set up can be found in the fluid limit model literature; see, e.g., [2]. The key concepts concern the tightness and the measures of stochastic processes, which leads to the convergence of the subsequences of the scaled processes.

Definition 2.1. *An absolutely continuous deterministic process* $\{(\boldsymbol{q}_t, \boldsymbol{\delta}_t) \in \mathbb{R}^{\mathbb{J}} \times \mathbb{R}^{\mathbb{I}} : t \in \mathbb{R}_+\}$ *is called a* (IQS) fluid model *with initial state* $\boldsymbol{q}_0 \in \mathbb{R}_+^{\mathbb{J}}$ *and arrival rates* $\boldsymbol{\lambda} \in [0, 1]^{\mathbb{J}}$ *if the following conditions hold:* ***(FM1)*** $\boldsymbol{q}_t = \boldsymbol{q}_0 + \boldsymbol{\lambda} t - \boldsymbol{\delta}_t \boldsymbol{A}$ *for* $t \in \mathbb{R}_+$*;* ***(FM2)*** $\boldsymbol{q}_t \ge 0$ *for* $t \in \mathbb{R}_+$*;* ***(FM3)*** $\sum_{\boldsymbol{s}\in\mathbb{I}} \delta_t(\boldsymbol{s}) = t$ *(i.e.,* $\|\boldsymbol{\delta}_t\|_1 = t$*) and* $\boldsymbol{\delta}_t \ge \boldsymbol{0}$ *for* $t \in \mathbb{R}_+$*;* ***(FM4)*** *For any* $\boldsymbol{s} \in \mathbb{I}$*,* $\delta_t(\boldsymbol{s})$ *is non-decreasing w.r.t.* t*. Consider another deterministic process* $\{\boldsymbol{\mu}_t \in \mathbb{R}_+ : t \in \mathbb{R}_+\}$*, which is called*

MAMA 2020 Boston, MA USA

an (fluid-level) admissible policy *for the IQS iff there exists a fluid model* $(\boldsymbol{q}_t, \boldsymbol{\delta}_t)$ *s.t.* $\boldsymbol{\mu}_t = \dot{\boldsymbol{\delta}}_t \forall t \in \mathbb{R}_+$ *at which* $\dot{\boldsymbol{\delta}}_t$ *exists.*

PROPOSITION 2.1. *Fix* $\boldsymbol{q} \in \mathbb{R}_+^{\mathbb{J}}$ *and* $\boldsymbol{\lambda} \in [0,1]^{\mathbb{J}}$. *Let* $\{\boldsymbol{\mu}_t \in \mathbb{R}_+^{\mathbb{I}} : t \in \mathbb{R}_+\}$ *be an integrable deterministic process and* $\{\boldsymbol{q}_t \in \mathbb{R}^{\mathbb{J}} : t \in \mathbb{R}_+\}$ *a process satisfying* $\dot{\boldsymbol{q}}_t = \boldsymbol{\lambda} - \boldsymbol{\mu}_t \boldsymbol{A}$ *with initial state* $\boldsymbol{q}_0$. *Then, the following statements are equivalent:* ***(AP1)*** $\boldsymbol{\mu}_t$ *is a fluid-level admissible policy;* ***(AP2)*** $\|\boldsymbol{\mu}_t\|_1 = 1$ *and* $\boldsymbol{q}_t \geq 0 \forall t \in \mathbb{R}_+$*;* ***(AP3)*** $\|\boldsymbol{\mu}_t\|_1 = 1$ *and* $\boldsymbol{\mu}_t \in \mathbb{U}(\boldsymbol{q}_t) \forall t \in \mathbb{R}_+$, *where* $\mathbb{U}(\boldsymbol{q}) := \{\boldsymbol{\mu} \in [0,1]^{\mathbb{I}} : (\mu A)(\boldsymbol{\rho}) \leq \lambda(\boldsymbol{\rho})$ *if* $q(\boldsymbol{\rho}) = 0\}$. *In this case,* $(\boldsymbol{q}_t, \boldsymbol{\delta}_t := \int_0^t \boldsymbol{\mu}_{t'} dt')$ *is the fluid model associated with fluid-level admissible policy* $\boldsymbol{\mu}_t$.

We next introduce a family of scaled processes and demonstrate that converging subsequences have fluid models as their limits, motivating our optimal control problems in §2.3.

Fix index $r \in \mathbb{Z}^+$ and let $\{(\boldsymbol{\mathcal{Q}}_t^r, \boldsymbol{\mathcal{A}}_t^r, \boldsymbol{\mathcal{D}}_t^r) : t \in \mathbb{Z}_+\}$ be a discrete-time stochastic model with initial state $\boldsymbol{\mathcal{Q}}^r$ as described in §2.1. We extend this discrete-time process to a continuous-time process by defining $\boldsymbol{\mathcal{A}}_t^r := (t - \lfloor t \rfloor)(\boldsymbol{\mathcal{A}}_{\lfloor t \rfloor + 1}^r - \boldsymbol{\mathcal{A}}_{\lfloor t \rfloor}^r) + \boldsymbol{\mathcal{A}}_{\lfloor t \rfloor}^r$, $\boldsymbol{\mathcal{D}}_t^r := (t - \lfloor t \rfloor)(\boldsymbol{\mathcal{D}}_{\lfloor t \rfloor + 1}^r - \boldsymbol{\mathcal{D}}_{\lfloor t \rfloor}^r) + \boldsymbol{\mathcal{D}}_{\lfloor t \rfloor}^r$, $\boldsymbol{\mathcal{Q}}_t^r := (t - \lfloor t \rfloor)(\boldsymbol{\mathcal{Q}}_{\lfloor t \rfloor + 1}^r - \boldsymbol{\mathcal{Q}}_{\lfloor t \rfloor}^r) + \boldsymbol{\mathcal{Q}}_{\lfloor t \rfloor}^r = \boldsymbol{\mathcal{Q}}^r + \boldsymbol{\mathcal{A}}_t^r - \boldsymbol{\mathcal{D}}_t^r \boldsymbol{A}$, where $\lfloor t \rfloor$ is the largest integer less than or equal to t.

For randomness $\boldsymbol{\omega}$, the scaled r-th system is defined by $(\hat{\boldsymbol{\mathcal{Q}}}_t^r(\boldsymbol{\omega}), \hat{\boldsymbol{\mathcal{A}}}_t^r(\boldsymbol{\omega}), \hat{\boldsymbol{\mathcal{D}}}_t^r(\boldsymbol{\omega})) := (r^{-1}\boldsymbol{\mathcal{Q}}_{rt}^r(\boldsymbol{\omega}), r^{-1}\boldsymbol{\mathcal{A}}_{rt}^r(\boldsymbol{\omega}), r^{-1}\boldsymbol{\mathcal{D}}_{rt}^r(\boldsymbol{\omega}))$. We assume that the initial state of the r-th system satisfies $r^{-1}\boldsymbol{\mathcal{Q}}_0^r \Rightarrow \boldsymbol{q}_0$, as $r \to \infty$, for some point $\boldsymbol{q}_0 \in \mathbb{R}_+^{\mathbb{J}}$, where the convergence is understood to be convergence in distribution.

For a fixed sample path $\boldsymbol{\omega}$, from the equations for $\|\boldsymbol{\mathcal{D}}_t\|_1$ and $(\boldsymbol{\mathcal{Q}}_t^r, \boldsymbol{\mathcal{A}}_t^r, \boldsymbol{\mathcal{D}}_t^r)$, we have $\hat{\mathcal{D}}_0(\boldsymbol{\rho}; \boldsymbol{\omega}) = 0$ and $\hat{\mathcal{D}}_t(\boldsymbol{\rho}; \boldsymbol{\omega}) \leq \|\hat{\boldsymbol{\mathcal{D}}}_t(\boldsymbol{\omega})\|_1 = t$ so that $\hat{\mathcal{D}}_t^r(\boldsymbol{\rho}; \boldsymbol{\omega}) - \hat{\mathcal{D}}_{t'}^r(\boldsymbol{\rho}; \boldsymbol{\omega}) \leq (t - t')$ for any $r > 0$ and $t \geq t' \geq 0$, which implies the tightness of $\hat{\mathcal{D}}_t^r$.

From the functional strong law of large numbers, we have $\lim_{r\to\infty} \sup_{0 \leq t \leq T} |\hat{\mathcal{A}}_t^r(\boldsymbol{\rho}; \boldsymbol{\omega}) - \lambda(\boldsymbol{\rho})t| = 0$ a.s. Hence, for any sequence $\{r_k\}$ with $\lim_{k\to\infty} r_k = \infty$, there exists a subsequence $\{r_{k_l}\}$ and absolutely continuous deterministic process $(\boldsymbol{q}_t, \boldsymbol{\delta}_t)$, which is a fluid model in Definition 2.1, s.t. $(\hat{\boldsymbol{\mathcal{Q}}}_t^{r_{k_l}}(\boldsymbol{\omega}), \hat{\boldsymbol{\mathcal{D}}}_t^{r_{k_l}}(\boldsymbol{\omega})) \to (\boldsymbol{q}_t, \boldsymbol{\delta}_t)$, a.s. uniformly on all compact sets as $l \to \infty$.

2.3 Fluid Model Optimal Control Problems

Define the total discounted delay cost over the entire time horizon under a fluid-level admissible policy $\{\boldsymbol{\mu}_t : t \in \mathbb{R}_+\}$ with initial state $\boldsymbol{q}_0$ by $c(\boldsymbol{\mu}_t; \boldsymbol{q}_0) := \int_0^\infty e^{-\beta t} \boldsymbol{c} \cdot \boldsymbol{q}_t dt$, where $\boldsymbol{q}_t$ is the deterministic function defined in **(FM1)** with $\boldsymbol{\delta}_t := \int_0^t \boldsymbol{\mu}_s ds$, β is the discount factor, and $\boldsymbol{c} \in (\mathbb{R}^+)^{\mathbb{J}}$ is the vector of cost coefficients. From **(AP2)** in Proposition 2.1, we then formulate the control problem **(OC)** as minimizing $c(\boldsymbol{\mu}_t; \boldsymbol{q}_0)$ over all admissible policies $\{\boldsymbol{\mu}_t : t \in \mathbb{R}_+\}$ subject to $\dot{\boldsymbol{q}}_t = \boldsymbol{\lambda} - \boldsymbol{\mu}_t \boldsymbol{A}$, $\boldsymbol{q}_t \geq \boldsymbol{0}$ and $\boldsymbol{\mu}_t \in \mathbb{U}$, $\forall t \in \mathbb{R}_+$, where $\mathbb{U} = \{\boldsymbol{\mu} \in [0,1]^{\mathbb{I}} : \|\boldsymbol{\mu}\|_1 = 1\}$ and the initial state of $\boldsymbol{q}_t$ is $\boldsymbol{q}_0$.

From Pontryagin's maximum principle [4] under appropriate conditions, the following proposition provides sufficient conditions for an optimal solution of the control problem.

PROPOSITION 2.2. *Let* $\boldsymbol{q}_0$ *be the initial condition of a fluid model. Let* $\{\boldsymbol{\mu}_t^* \in \mathbb{R}_+^{\mathbb{I}} : t \in \mathbb{R}_+\}$ *be a fluid-level admissible policy, and let* $\boldsymbol{q}_t^* = \boldsymbol{q}_0 + \boldsymbol{\lambda} t + \int_0^t \boldsymbol{\mu}_{t'}^* \boldsymbol{A} dt'$ *be the associated QL process. Assume there exists a continuous process* $\{\boldsymbol{p}_t \in \mathbb{R}^{\mathbb{J}} : t \in \mathbb{R}_+\}$ *with piecewise continuous* $\dot{\boldsymbol{p}}_t$ *and a process* $\{\boldsymbol{\eta}_t \in \mathbb{R}_+^{\mathbb{J}} : t \in \mathbb{R}_+\}$ *s.t. the following conditions are satisfied:* ***(C1)*** $\boldsymbol{\mu}_t^* \in \arg\max\{\mu A p_t : \boldsymbol{\mu} \in \mathbb{U}\}$*;* ***(C2)*** $\dot{\boldsymbol{p}}_t - \beta \boldsymbol{p}_t = \boldsymbol{c} - \boldsymbol{\eta}_t$*;* ***(C3)*** $\boldsymbol{q}_t^* \cdot \boldsymbol{\eta}_t = 0$, $\boldsymbol{q}_t^* \geq 0$, $\boldsymbol{\eta}_t \geq 0$*;* ***(C4)*** $\liminf_{t\to\infty} \boldsymbol{p}_t \cdot (\boldsymbol{q}_t^* - \boldsymbol{q}_t) \geq 0$ *for any fluid model* $(\boldsymbol{q}_t, \boldsymbol{\delta}_t)$ *with initial condition* $\boldsymbol{q}_0$. *Then,* $\{\boldsymbol{\mu}_t^* : t \in \mathbb{R}_+\}$ *is an optimal solution to the optimal control problem* ***(OC)***.

Algorithm 1 Find critical threshold at state $\boldsymbol{q}$ in $\mathbb{W}$

Input: None, **Output:** An integer

Set $l = 1$ and $h = \min\{k : \exists \boldsymbol{s} \in \mathbb{J}$ s.t. $w(\boldsymbol{s}) = \tau_k,\ q_{\boldsymbol{\rho}} \neq 0\ \forall \boldsymbol{\rho} \in \boldsymbol{s}\}$
Solve Problem $(P_{\boldsymbol{q},\tau})$ with $\tau = \tau_l$, obtain an optimal value γ_l and an optimal solution $\boldsymbol{\nu}^*$
if $\mathbb{Q}(\boldsymbol{q}, \tau_l, \gamma_l) \neq \emptyset$ **then**
    return l
Solve Problem $(P_{\boldsymbol{q},\tau})$ with $\tau = \tau_h$, obtain an optimal value γ_h and an optimal solution $\boldsymbol{\nu}^*$
if $\mathbb{Q}(\boldsymbol{q}, \tau_h, \gamma_h) \neq \emptyset$ **then**
    return h
while $l < h - 1$ **do**
    Set $m = \lfloor \frac{l+h}{2} \rfloor$ and $\tau = \tau_m$
    Solve Problem $(P_{\boldsymbol{q},\tau})$ with $\tau = \tau_m$, obtain an optimal value γ_m and an optimal solution $\boldsymbol{\nu}^*$
    if $\mathbb{Q}(\boldsymbol{q}, \tau_m, \gamma_m) \neq \emptyset$ **then**
        return m
    else
        if $\|\boldsymbol{\nu}^*\|_1 > 1$ **then**
            Set $h = m$
        else
            Set $l = m$
return $-l$

3. OPTIMAL CONTROL

In this section, we present and analyze algorithms that render the optimal fluid-cost scheduling policy, i.e., the optimal solution to the control problem **(OC)** of §2.3. We refer to the stochastic model in §2.1 as the pre-limit model and refer to fluid model in §2.2 as the limit system. For the pre-limit model, recall that a basic schedule is a collection of queues from each of which a packet can depart simultaneously, where $\mathbb{J} := [n] \times [n]$ denotes the set of queues. A basic schedule is represented by a $|\mathbb{J}|$-dimensional binary vector $\boldsymbol{s} = [s(\boldsymbol{\rho}) \in \{0,1\} : \boldsymbol{\rho} \in \mathbb{J}]$, where $s(\boldsymbol{\rho}) = 1$ iff $\boldsymbol{\rho}$ is in the collection composing the basic schedule. We use $\boldsymbol{\rho} \in \boldsymbol{s}$ if $s(\boldsymbol{\rho}) = 1$, for $\boldsymbol{\rho} \in \mathbb{J}$ and $\boldsymbol{s} \in \mathbb{I}$. For a basic schedule $\boldsymbol{s} \in \mathbb{I}$, with $\mathbb{I}$ the set of all basic schedules given above, we define the *weight* of $\boldsymbol{s}$ by $w(\boldsymbol{s}) := \sum_{\boldsymbol{\rho} \in \boldsymbol{s}} c(\boldsymbol{\rho})$, where $\boldsymbol{c} \in (\mathbb{R}^+)^{\mathbb{J}}$ is the cost coefficient vector introduced in **(OC)**.

While time in the pre-limit system is discrete with QL vector $\boldsymbol{\mathcal{Q}}_t \in \mathbb{Z}_+^{\mathbb{J}}$ at time $t \in \mathbb{Z}_+$, time in the limit system is continuous with the state space of (fluid) QL vectors $\boldsymbol{q}_t$ given by $\mathbb{R}_+^{\mathbb{J}}$. From Proposition 2.1, we define a *(fluid-level) schedule* by a convex combination of basic schedules and represent it as an $|\mathbb{I}|$-dimensional vector $\boldsymbol{\mu} = [\mu(\boldsymbol{s}) \in [0,1] : \boldsymbol{s} \in \mathbb{I}]$ with $\|\boldsymbol{\mu}\|_1 = 1$, where $\mu(\boldsymbol{s})$ is the coefficient of schedule $\boldsymbol{s}$. Furthermore, schedule $\boldsymbol{\mu}$ is *admissible* at state $\boldsymbol{q} \in \mathbb{R}_+^{\mathbb{J}}$ iff $\boldsymbol{\mu} \in \mathbb{U}(\boldsymbol{q})$, as defined in **(AP3)**.

3.1 Critical Thresholds

We now introduce, for each state $\boldsymbol{q} \in \mathbb{R}_+^{\mathbb{J}}$, a family of LP problems, indexed by non-negative real numbers, from which we construct an (admissible) schedule associated with the LP. These schedules are instrumental to the development of the optimal control algorithms in Sec. 3.2. For a given state $\boldsymbol{q}$ and a real value $\tau \in \mathbb{R}_+$, define sets $\mathbb{I}_\tau \subset \mathbb{I}$ and $\mathbb{J}_{\boldsymbol{q}} \subset \mathbb{J}$

by $\mathbb{I}_\tau := \{\boldsymbol{s} \in \mathbb{I} : w(\boldsymbol{s}) \geq \tau\}$, $\mathbb{J}_{\boldsymbol{q}} := \{\boldsymbol{\rho} \in \mathbb{J} : q(\boldsymbol{\rho}) = 0\}$, respectively, and define $\boldsymbol{w}_\tau := [w(\boldsymbol{s}) - \tau : \boldsymbol{s} \in \mathbb{I}_\tau] \in \mathbb{R}_+^{\mathbb{I}_\tau}$. For τ with $\mathbb{I}_\tau \neq \emptyset$, we formulate the following LP problem:

$$\max \quad \boldsymbol{w}_\tau \cdot \boldsymbol{\nu}, \quad \text{s.t.} \quad \boldsymbol{\nu} \boldsymbol{A}_{\tau,\boldsymbol{q}} \leq \boldsymbol{\lambda}_{\boldsymbol{q}}, \quad \boldsymbol{\nu} \geq \boldsymbol{0}, \qquad (P_{\boldsymbol{q},\tau})$$

where $\boldsymbol{A}_{\tau,\boldsymbol{q}} := [A(\boldsymbol{s},\boldsymbol{\rho}) : \boldsymbol{s} \in \mathbb{I}_\tau, \boldsymbol{\rho} \in \mathbb{J}_{\boldsymbol{q}}] \in \{0,1\}^{\mathbb{I}_\tau \times \mathbb{J}_{\boldsymbol{q}}}$, $\boldsymbol{\lambda}_{\boldsymbol{q}} := [\lambda(\boldsymbol{\rho}) : \boldsymbol{\rho} \in \mathbb{J}_{\boldsymbol{q}}] \in [0,1]^{\mathbb{J}_{\boldsymbol{q}}}$, $\boldsymbol{\nu} \in \mathbb{R}^{\mathbb{I}_\tau}$ is the vector of decision variables. Note: if $\tau = 0$, then $\mathbb{I}_0 = \mathbb{I}$, $\boldsymbol{w}_0 = \boldsymbol{A}\boldsymbol{c}$.

THEOREM 3.1. *For any state $\boldsymbol{q}$, there exists a $\tau = \tau(\boldsymbol{q}) \in \mathbb{R}_+$ s.t. Problem $(P_{\boldsymbol{q},\tau})$ has an optimal solution $\boldsymbol{\nu}$ that can be extended to an admissible schedule at state $\boldsymbol{q}$; namely, $\|\boldsymbol{\nu}\|_1 = 1$. We call such τ a* critical threshold *of state $\boldsymbol{q}$.*

Algorithm 2 Find critical threshold at state $\boldsymbol{q}$ in (τ_{l+1}, τ_l)

Input: integer l s.t.: 1-norm of any optimal solution to Problem $(P'_{\boldsymbol{q},\tau})$ with $\tau = \tau_l$ is less than 1; and 1-norm of any optimal solution to Problem $(P'_{\boldsymbol{q},\tau})$ with $\tau = \tau_{l+1}$ is greater than 1
Output: a critical threshold $\tau \in (\tau_{l+1}, \tau_l)$

Set $\bar{\boldsymbol{w}} = [w(\boldsymbol{s}) : \boldsymbol{s} \in \mathbb{I}_{\tau_l}]$, and $k = 0$
Set $\tau_0^L = \tau_l$ and obtain a basic optimal solution $\boldsymbol{\nu}_0^L$ to Problem $(P'_{\boldsymbol{q},\tau})$ with $\tau = \tau_0^L$
Set $\tau_0^S = \tau_{l+1}$ and obtain a basic optimal solution $\boldsymbol{\nu}_0^S$ to Problem $(P'_{\boldsymbol{q},\tau})$ with $\tau = \tau_0^S$
while True **do**
  Set $\tau_k^M := \frac{\bar{\boldsymbol{w}} \cdot (\boldsymbol{\nu}_k^S - \boldsymbol{\nu}_k^L)}{\|\boldsymbol{\nu}_k^S\|_1 - \|\boldsymbol{\nu}_k^L\|_1}$
  Solve Problem $(P'_{\boldsymbol{q},\tau})$ with $\tau = \tau_k^M$, obtain optimal value γ^* and basic optimal solution $\boldsymbol{\nu}_k^M$
  if $\mathbb{Q}'(\boldsymbol{q}, \tau_k^M, \gamma^*) \neq \emptyset$ **then**
    return τ_k^M
  else
    if $\|\boldsymbol{\nu}_k^M\|_1 > 1$ **then**
      Set $(\tau_{k+1}^S, \boldsymbol{\nu}_{k+1}^S) = (\tau_k^M, \boldsymbol{\nu}_k^M)$ and $(\tau_{k+1}^L, \boldsymbol{\nu}_{k+1}^L) = (\tau_k^L, \boldsymbol{\nu}_k^L)$
    else
      Set $(\tau_{k+1}^L, \boldsymbol{\nu}_{k+1}^L) = (\tau_k^M, \boldsymbol{\nu}_k^M)$ and $(\tau_{k+1}^S, \boldsymbol{\nu}_{k+1}^S) = (\tau_k^S, \boldsymbol{\nu}_k^S)$
    Set $k = k + 1$

Our solution to find the critical threshold in Theorem 3.1 is presented in Algorithms 1 and 2. Letting γ denote the optimal value of Problem $(P_{\boldsymbol{q},\tau})$, then τ is a critical threshold at state $\boldsymbol{q}$ iff the following set is nonempty:

$$\mathbb{Q}(\boldsymbol{q}, \tau, \gamma) := \{\boldsymbol{\nu} : \boldsymbol{w}_\tau \cdot \boldsymbol{\nu} = \gamma, \|\boldsymbol{\nu}\|_1 = 1, \boldsymbol{\nu} \boldsymbol{A}_{\tau,\boldsymbol{q}} \leq \boldsymbol{\lambda}_{\boldsymbol{q}}\}. \quad (1)$$

Note that all constraints in (1) are linear, and thus $\mathbb{Q}(\boldsymbol{q}, \tau, \gamma)$ is a polyhedron, which implies that the emptiness of $\mathbb{Q}(\boldsymbol{q}, \tau, \gamma)$ can be checked quickly through the solution of an LP.

Define $\mathbb{W} := \{w(\boldsymbol{s}) : \boldsymbol{s} \in \mathbb{I}\} = \{\tau_1, \tau_2, \dots\}$ to be the ordered set of all (distinct) weights of schedules in $\mathbb{J}$ with $\tau_i > \tau_{i+1}$ for $i = 1, 2, \dots$. If $\mathbb{W}$ contains a critical threshold, Algorithm 1 returns a positive integer m s.t. τ_m is a critical threshold, which is the key element needed for our optimal control policy. Otherwise, Algorithm 1 returns $-l$ for some $l \in \mathbb{Z}^+$; and if a critical threshold exists in $\mathbb{R}_+$ (but not in $\mathbb{W}$), then it is between τ_{l+1} and τ_l. We define $\bar{\boldsymbol{w}} := [w(\boldsymbol{s}) : \boldsymbol{s} \in \mathbb{I}_{\tau_l}]$ and formulate another LP for $\tau \in (\tau_{l+1}, \tau_l)$:

$$\max \ \bar{\boldsymbol{w}} \cdot \boldsymbol{\nu} - \tau \|\boldsymbol{\nu}\|_1, \quad \text{s.t.} \ \boldsymbol{\nu} \boldsymbol{A}_{\tau_l,\boldsymbol{q}} \leq \boldsymbol{\lambda}_{\boldsymbol{q}}, \boldsymbol{\nu} \geq \boldsymbol{0}. \qquad (P'_{\boldsymbol{q},\tau})$$

Algorithm 2 finds a critical threshold in (τ_{l+1}, τ_l).

3.2 Optimal Control Algorithm

By exploiting the critical threshold for any state $\boldsymbol{q}$ from the previous section, we now introduce an optimal control algorithm and show that it renders an optimal solution to the optimal control problem **(OC)**.

Algorithm 3 Optimal Control Algorithm for initial state $\boldsymbol{q}_{t=0}$

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

PROPOSITION 3.1. *In Algorithm 3, we have that $\boldsymbol{\mu}_t^*$ is a fluid-level admissible policy and $\boldsymbol{q}_t^*$ is the continuous process satisfying $\dot{\boldsymbol{q}}_t^* = \boldsymbol{\lambda} - \boldsymbol{\mu}_t^* \boldsymbol{A}$ with initial state $\boldsymbol{q}_{t=0}$.*

THEOREM 3.2. *Assume that for arrival rate vector $\boldsymbol{\lambda}$, $(\boldsymbol{q}_t^*, \boldsymbol{\mu}_t^*)$ be an admissible pair under Algorithm 3 which empties the system in finite time. Then, $(\boldsymbol{q}_t^*, \boldsymbol{\mu}_t^*)$ is an optimal solution to problem **(OC)**.*

3.3 Relationship with cμ Policy

Given an arrival rate vector $\boldsymbol{\lambda}$ and initial QL $\boldsymbol{q}_0$ s.t. $\lambda(i,j) = q_0(i,j) = 0, \forall i \in [n]$ and $j \in [n] \backslash \{1\}$, the $n \times n$ IQS is equivalent to n parallel queues with one server. The cμ-policy is well-known for this case to be an optimal policy that minimizes the discounted total cost over an infinite horizon in both the stochastic and fluid models (see [1]); and, in this case, Algorithm 3 follows the cμ-policy in the fluid model.

However, the cμ-policy is not optimal for the $n \times n$ IQS in general. In fact, even a fluid model under the cμ-policy can be unstable. Using a formal definition of stability (*weak stability*) for fluid models, we provide an example in [3] showing that the cμ-policy is not weakly stable. In contrast, we have the following proposition for Algorithm 3.

PROPOSITION 3.2. *Algorithm 3 is weakly stable for every arrival rate $\boldsymbol{\lambda}$ s.t. $\sum_{k=1}^n \lambda(k,j) < 1$ and $\sum_{k=1}^n \lambda(i,k) < 1$ for all $i, j \in [n]$.*

4. REFERENCES

[1] N. Bäuerle, U. Rieder. Optimal control of single-server fluid networks. *QUESTA*, 35(1):185–200, Jul 2000.

[2] J. G. Dai. On positive harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. *Ann. Appl. Probab.*, 5(1):49–77, 02 1995.

[3] Y. Lu, M.S. Squillante, T. Suk. Optimal delay-cost scheduling control in fluid models of general $n \times n$ input-queued switches. arXiv:1910.13657, 2019.

[4] L. Pontryagin, V. Boltyanskii, R. Gamkrelidze, E. Mishchenko. *The Mathematical Theory of Optimal Processes.* Interscience, New York, 1962.