Competitive Online Path-Aware Path Selection

Ying Cao, Siyuan Yu, Xiaoqi Tan, Danny H.K. Tsang

Abstract

This paper studies an online path selection problem and proposes online mechanisms for a network operator to sequentially update link prices. The aim is to incentivize online-arriving agents to join the network and select paths in a manner that maximizes the social welfare, which comprises both system profit and the quality of service experienced by agents. Competitive analysis is adopted to analyze the performance of the proposed online mechanism, whose best achievable competitive ratio is 4. Sufficient and necessary conditions on a competitive mechanism are established. Moreover, the performance limit of the celebrated multiple-the-index pricing scheme is also analyzed.

1. Introduction

End-host path selection has become a critical component of network performance and reliability in future Internet architectures [1, 2]. While traditional path selection algorithms prioritize the path with the lowest latency or congestion, research in the past two decades has shown the benefits of path-aware path selection. This approach takes into account other factors such as the path's available bandwidth, traffic types, and QoS requirements. However, to the best of our knowledge, theoretical research in this field has primarily focused on the long-term equilibrium that is achieved in dynamic rate evolutions [2], while the transient performance when the system starts is understudied. A well-designed path selection policy is especially important at the beginning of a network session, particularly when users remain in the system for an extended period. This is because an effective initial path selection can eliminate the need for frequent path switching, which can incur additional switching costs in dynamic path selection policies and is undesirable according to the IETF [3]. By reducing the frequency of path switching, a good path selection policy can help improve network performance and reduce the likelihood of congestion or packet loss. This gap highlights the need for a better understanding of how path-aware path selection algorithms should be designed during the system's initial stage.

In this work, we study the following transient scenario when users/agents come sequentially to nodes/hosts in a zero-load network. Through the investi-

Email addresses: ycaoan@connect.ust.hk (Ying Cao), syu3@ualberta.ca (Siyuan Yu), xiaoqi.tan@ualberta.ca (Xiaoqi Tan), eetsang@ust.hk (Danny H.K. Tsang)

gation of this scenario with the assumption that future agents will arrive in an adversarial manner, the objective is to enhance the network's ability to handle the uncertainty of future network dynamics during the initial stage. We assume agents are *strategic*, in that each of them holds an arbitrary valuation of allocated resources, and may misreport her valuation to the network for her own benefit. The valuation can be an indicator of the agent's type. By considering strategic behavior in the selection of paths, we aim to design path selection mechanisms that are more efficient and effectively cater to the requirements of both the network and its users. The network operator sets unit price of network resources based on the network state such as available link/edge capacity and path congestion, and the agent decides whether to join the network by comparing her valuation with the lowest price among all feasible paths. After joining the network, agents pay the price to the network operator for the network resources consumed when their data are routed through the path selected. The goal of the network operator is to maximize the social welfare, i.e., the total valuation of agents in the network minus the dissatisfaction (e.g., packet loss or delay experienced) caused by poor network states. In this work, the network operator achieves the goal by sequentially setting proper prices for links as an incentive for agents to use the network resources in a way that induces a better network state, for example, a more balanced network load. In economic mechanism design, such sequential pricing schemes are known as *posted-price mechanisms*, which enjoy the desired property of being incentive-compatible, i.e., agents are willing to report their true valuation.

Mechanism design for communication networks in the offline setting has been studied in [4]. Price-of-anarchy (PoA), the ratio between the cost of an equilibrium under decentralized decision making and that from the centralized cost minimization, was the performance metric there. When agents arrive online, a unified algorithmic framework for online resource allocation was proposed in [5]. Only the allocation of a single resource was considered, and the allocation of multiple resources is notably harder [6]. Combinatorial auctions were studied in [7], in which resources were packed into different bundles and priced bundle-wise. Agents bid for bundles and consume the resources in the bundle won. Additional resources are produced with polynomial or linear cost functions.

Online routing, as an independent problem, has been garnering research interest from theoreticians since 90s. Since then, routing problems have been conventionally classified into cost minimization and benefit maximization problems [8]. However, none of previous works considered the problem of maximizing benefit while regulating the network state (in the form of link costs) as in this work.

(**Our contributions**) We study how to design competitive online mechanisms for path-aware path selection. In this work, the impact of congested links on agents' service degradation is reflected by the cost in proportion to the total number of packets in the network. The cost grows to infinity and calls on the use of the effective capacity concept newly introduced in this work. We derive the sufficient and necessary conditions on the pricing mechanism to be competitive, and show that these conditions relate the competitiveness of the proposed online mechanism to the existence of solutions to a differential equation with two boundary conditions. The main technical contribution is to identify the boundary conditions, especially the one on the right hand side, which relies on finding the worst-case scenario. Apart from developing general conditions on competitive online mechanisms, the performance of a celebrated pricing scheme, multiple-the-index pricing [9], is studied and shown inferior to our design.

2. Problem Statement

Consider a packet-switched network that consists of E edges, where edge e is endowed with capacity c_e and the collection of all edges is denoted as $\mathcal{E}$. Agents are interested in transferring data between a source node and a destination node. They arrive at the network one by one and request to route data through a path that connects the source and destination nodes. There are N agents in total, but the network operator does not know the value of N or the information about future agents throughout time. Despite the lack of future information, the network operator sets price for each link at the beginning and may update the price according to the varying network state and the sequential arrival of agents.

The ith agent carries the following information: the private valuation v_i, the average data rate r_i, the source and destination node pair (s_i, t_i). Upon the arrival of agent i, she decides whether to join the network or not based on her private valuation and the price charged by the network operator; if agent i decides to join, she continues to decide how to route her data through the network and generates packet flows along the route; otherwise, agent i leaves the network. The packet flows follow a Poisson process with mean rate r_i. Given s_i and t_i, there is a set of possible paths from s_i to t_i denoted as $\mathcal{O}_i$.

The routing decision of agent i is denoted as $x_i \in \{0,1\}^{|\mathcal{O}_i|}$; if $x_{ij} = 1$, then the jth path is chosen to route the flow of agent i; if $\sum_{j\in\mathcal{O}_i} x_{ij} = 0$, it represents that the agent chooses to leave the system. The sum of mean rates for all flows passing through edge e is $\lambda_e = \sum_{i\in[N]} r_i \sum_{j\in\mathcal{O}_i} \delta_{ij}^e x_{ij}$, where $\delta_{ij}^e = 1$ indicates that edge e is on the jth path in $\mathcal{O}_i$, and $\delta_{ij}^e = 0$ indicates the opposite. In this work, each edge is modeled as an M/M/1 queue with arrival rate λ_e and service rate c_e, and the service quality degradation is quantified by the total number of packets in the network $\sum_{e\in\mathcal{E}} f(\rho_e)$, where

$$f(\rho_e) = \begin{cases} \frac{\rho_e}{1-\rho_e}, & 0 \le \rho_e < 1, \\ \infty, & \rho_e \ge 1, \end{cases}$$

and $\rho_e = \frac{\lambda_e}{c_e}$ is the utilization of link e. It is well-accepted in queuing theory that the number of packets in the network is positively related to the average network delay and harms the service quality to some extent. Other typical network state preferences, such as maximizing the minimum load, can also be incorporated by enforcing different f's. In this regard, f can be viewed as a regularizer of the network state.

The utility of agent i is defined as her valuation minus the payment to the network, and the utility of the network operator is defined as the total payment collected from agents minus the total service quality degradation as a result of link congestion. The social welfare of the network is defined as the sum of agent utility and the network utility, and the payments are cancelled out. If uncertainties about all future agents are resolved, i.e., the private and public information of all agents $\{v_i, r_i, \mathcal{O}_i\}_{i\in[N]}$ is disclosed at the beginning, to maximize the social welfare, the network operator just needs to solve the following optimization problem:

$$\max_{x,\rho} \quad \sum_{i\in[N]} v_i \sum_{j\in\mathcal{O}_i} x_{ij} - \gamma \sum_{e\in\mathcal{E}} f(\rho_e) \tag{1}$$

$$s.t. \quad \rho_e c_e = \sum_{i\in[N]} r_i \sum_{j\in\mathcal{O}_i} \delta_{ij}^e x_{ij}, \forall e \in \mathcal{E}, \qquad (p_e)$$

$$\sum_{j\in\mathcal{O}_i} x_{ij} \le 1, \forall i \in [N], \qquad (\mu_i)$$

$$x_{ij} \in \{0,1\}, \forall i \in [N], j \in \mathcal{O}_i, \tag{2}$$

where γ is a tradeoff parameter between the total agent valuation and the congestion effect. The sequence of all agents is called an instance. In the following, without loss of generality, we set $\gamma = 1$. It is worth noting that, although payments are cancelled out in the social welfare, the payment rule is the key to an effective pricing mechanism. We will show the important role it plays in influencing agent decisions and quantify the requirement of a competitive payment rule in Section 3.

Problem (1) is combinatorial and involves binary decision variables. Relaxing the binary constraints in Eq. (2) to fractional ones $x_{ij} \in [0,1]$ leads to a convex program whose dual problem is expressed as follows:

$$\min_{\mu,p} \quad \sum_{i\in[N]} \mu_i + \sum_{e\in\mathcal{E}} f^*(p_e c_e) \tag{3}$$

$$\text{s.t.} \quad \mu_i \ge v_i - r_i \sum_{e\in\mathcal{E}} \delta_{ij}^e p_e, \forall i \in [N], j \in \mathcal{O}_i,$$

where $f^*(y) = \sup_{\rho\in[0,\infty)}[y\rho - f(\rho)] = \begin{cases} (\sqrt{y}-1)^2, \text{if } y \ge 1, \\ 0, \text{if } y < 1. \end{cases}$

Both Problem (1) and Problem (3) cannot be solved in the traditional offline manner by the network operator because agent valuations are kept private and other agent information is not known until the arrival. Our key idea of dealing with this lack of information is twofold: (i) *we design a sequence of online dual variables* $\{p_e^{(i)}\}_{i\in\{0\}\bigcup[N]}$ *to approximate the optimal offline dual solution* p_e^* *to Problem* (3), and (ii) *let the online dual variables guide the design of feasible online primal decisions and then invoke weak duality to bound the primal objective.* The metric measuring the performance of online mechanisms is *competitive ratio*. An online mechanism is α-competitive if $\max_{\mathsf{I}} \frac{\mathsf{OPT(I)}}{\mathsf{ALG(I)}} \leq \alpha$, where $\mathsf{OPT(I)}$ and $\mathsf{ALG(I)}$ are the optimal objective value of Problem 1 and the objective value produced by the online algorithm when facing instance I, respectively. By definition, $\alpha \geq 1$ always holds, and the closer to 1 the better.

3. Major Results

Algorithm 1: Posted-Price Mechanism for Path Selection ($\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$)

Input: ϕ_e, $\rho_e^{(0)} = 0, \forall e$. Set $i = 1$.
while *a new agent i arrives* **do**
  if $\exists j \in \mathcal{O}_i, v_i \geq \sum_e \delta_{ij}^e r_i \phi_e(\rho_e^{(i-1)})$ **then**
    Agent i selects the min-cost path, i.e., set $x_{ij_i} = 1$ with j_i given by:

$$j_i = \arg\min_{j \in \mathcal{O}_i} \sum_e \delta_{ij}^e r_i \phi_e(\rho_e^{(i-1)}) \tag{4}$$

    If there is more than one solution to Eq. (4), choose one among them with equal probability;
    Update the link utilization: $\rho_e^{(i)} \leftarrow \rho_e^{(i-1)} + \frac{r_i}{c_e}, \forall e \in \{e | \delta_{ij_i}^e = 1\}$;
  else
    The agent leaves;
  end
  $i = i + 1$;
end

To facilitate the analysis, we make the following two assumptions [5, 6]. It is important to note that designing competitive algorithms for the general case is not feasible.

Assumption 1. *Agents are rational in the sense that their valuations are upper bounded*

$$v_i / r_i \leq \bar{p} \quad \textit{for all } i \in [N]. \tag{5}$$

We also assume that $\bar{p}$ is known to the network operator.

Assumption 2. *The rate r_i is much smaller than the capacity of any link, i.e., $r_i \ll c_e, \forall i \in [N], e \in \mathcal{E}$.*

The proposed algorithm $\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$ takes a carefully designed ϕ function as input, which is called a pricing scheme and essentially a mapping from edge utilization levels to edge prices. Whenever a new agent arrives, on each path she may choose, a new edge price is presented to her according to ϕ and the current edge utilization. The agent compares her valuation with the price of all possible paths and selects the one with the lowest price. The following two theorems (Theorem 1 and 2) relate the existence of a competitive pricing scheme with the existence of solutions to a differential equation with boundary conditions. Theorem 1 provides sufficient conditions on a differential equation for $\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$ (Eq. (6)) to be α-competitive.

Theorem 1 (Sufficiency)**.** *For any given $\alpha \geq 1$, $\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$ is α-competitive if $\boldsymbol{\phi} = (\phi_e)_{\forall e \in \mathcal{E}}$ and $\phi_e : [0, \bar{\rho}_e] \to \mathbb{R}$ is an analytic and non-decreasing solution to the following differential equation with boundary conditions:*

$$\begin{cases} \left(1 - (c_e \phi_e)^{-1/2}\right) \phi_e' = \alpha \left(\phi_e - f'/c_e\right), \\ \phi_e(0) = \frac{1}{c_e}, \phi_e(\bar{\rho}_e) \geq \bar{p}, \end{cases} \tag{6}$$

where $\bar{\rho}_e$ is such that $f'(\bar{\rho}_e) = \bar{p} c_e$.

Proof. According to the online primal-dual framework [10], there are three steps to deriving the sufficient conditions for an α-competitive algorithm:

1. Primal and dual initialization: P_0 and D_0 are the primal and dual objective values before any agent joins.

2. Primal and dual incremental inequality: $P_i - P_{i-1} \geq \frac{1}{\alpha}(D_i - D_{i-1}), \forall i \in [N]$. P_i and D_i are the primal and dual objective values after processing the ith agent, i.e., $P_i = \sum_{i' \leq i} v_{i'} \sum_{j \in \mathcal{O}_i} x_{i'j} - \sum_{e \in \mathcal{E}} f(\rho_e^{(i)})$ and $D_i = \sum_{i' \leq i} \mu_{i'} + \sum_{e \in \mathcal{E}} f^*(p_e^{(i)})$, where $\rho_e^{(i)}$ is the utilization level of link e after processing the ith agent, and $p_e^{(i)}$ are the dual variable updated after processing the ith agent.

3. Primal and dual feasibility: $\{x_i\}_{i \in [N]}$ are feasible for the primal problem, and $p_e^{(N)}$ is feasible for the dual problem.

By combining the above three steps, it is obvious that $P_N \geq \frac{1}{\alpha} D_N + P_0 - \frac{1}{\alpha} D_0 \geq \frac{1}{\alpha} OPT + P_0 - \frac{1}{\alpha} D_0$, which means an asymptotic competitive ratio of α. Next, we shall continue our analysis step by step.

First, before any agent arrives, the primal objective $P_0 = 0$, and the dual objective $D_0 = \sum_{e \in \mathcal{E}} f_e^*(p_e^{(0)} c_e) \geq 0$, where the equality holds when $p_e^{(0)} = \frac{1}{c_e}$. Let $p_e^{(i-1)} = \phi_e(\rho_e^{(i-1)})$ be a function of the utilization level $\rho_e^{(i-1)}$ and $\mu_i = \left(v_i - \sum_{e \in \mathcal{E}_i} r_i p_e^{(i-1)}\right)^+$. With $\phi_e(0) = \frac{1}{c_e}$, we have $P_0 = D_0$. Notice that the primal objective will increase only when an agent joins the network. If the ith agent leaves, the primal objective remains the same, and we keep the dual variables unchanged $p_e^{(i)} = p_e^{(i-1)}$, then the primal and dual incremental inequality is automatically satisfied for the ith agent. If the ith agent joins, $\sum_{j \in \mathcal{O}_i} x_{ij} = 1$. Define $\mathcal{E}_i$ as the edges that route the ith agent. The increment of the primal objective is

$$P_i - P_{i-1} = v_i - \sum_{e\in\mathcal{E}_i}\left[f(\rho_e^{(i)}) - f(\rho_e^{(i-1)})\right]$$
$$= \mu_i + \sum_{e\in\mathcal{E}_i} r_i p_e^{(i-1)} - \sum_{e\in\mathcal{E}_i}\left[f(\rho_e^{(i)}) - f(\rho_e^{(i-1)})\right], \tag{7}$$

where the last inequality follows from the way of selecting μ_i and p_e, $\mu_i = v_i - \sum_{e\in\mathcal{E}_i} r_i p_e^{(i-1)}$ because $v_i \geq \sum_{e\in\mathcal{E}_i} r_i p_e^{(i-1)}$ when the item i is admitted. Similarly, the increment of the dual objective is

$$D_i - D_{i-1} = \mu_i + \sum_{e\in\mathcal{E}_i}\left[f^*(p_e^{(i)}c_e) - f^*(p_e^{(i-1)}c_e)\right]. \tag{8}$$

The primal and dual incremental inequality is implied by the following inequality:

$$\sum_{e\in\mathcal{E}_i} r_i p_e^{(i-1)} - \sum_{e\in\mathcal{E}_i}\left[f(\rho_e^{(i)}) - f(\rho_e^{(i-1)})\right] + \left(1-\frac{1}{\alpha}\right)\mu_i \geq \frac{1}{\alpha}\sum_{e\in\mathcal{E}_i}\left[f^*(p_e^{(i)}c_e) - f^*(p_e^{(i-1)}c_e)\right]. \tag{9}$$

Because $\mu_i \geq 0$, the inequality above is implied by

$$\sum_{e\in\mathcal{E}_i} r_i p_e^{(i-1)} - \sum_{e\in\mathcal{E}_i}\left[f(\rho_e^{(i)}) - f(\rho_e^{(i-1)})\right] \geq \frac{1}{\alpha}\sum_{e\in\mathcal{E}_i}\left[f^*(p_e^{(i)}c_e) - f^*(p_e^{(i-1)}c_e)\right],$$

which is further implied by the individual inequalities over links:

$$r_i p_e^{(i-1)} - \left[f(\rho_e^{(i)}) - f(\rho_e^{(i-1)})\right] \geq \frac{1}{\alpha}\left[f^*(p_e^{(i)}c_e) - f^*(p_e^{(i-1)}c_e)\right]. \tag{10}$$

By dividing $r_i = c_e(\rho_e^{(i)} - \rho_e^{(i-1)})$ at both sides, we have $\forall e \in \mathcal{E}_i, \quad p_e^{(i-1)} - \frac{1}{c_e}\frac{f(\rho_e^{(i)})-f(\rho_e^{(i-1)})}{\rho_e^{(i)}-\rho_e^{(i-1)}} \geq \frac{1}{\alpha}\cdot\frac{f^*(p_e^{(i)}c_e)-f^*(p_e^{(i-1)}c_e)}{(p_e^{(i)}-p_e^{(i-1)})c_e}\cdot\frac{p_e^{(i)}-p_e^{(i-1)}}{\rho_e^{(i)}-\rho_e^{(i-1)}}$. By Assumption 2, it is equivalent to

$$\phi_e\left(\rho_e^{(i-1)}\right) - \frac{1}{c_e}f_e'\left(\rho_e^{(i-1)}\right) \geq \frac{1}{\alpha}f^{*'}(\phi_e(\rho_e^{(i-1)})c_e)\phi_e'(\rho_e^{(i-1)}),$$

which implies $P_i - P_{i-1} \geq \frac{1}{\alpha}(D_i - D_{i-1})$. Thus, a sufficient condition for the primal and dual incremental inequality is

$$\phi_e(\rho) - \frac{1}{c_e}f'(\rho) \geq \frac{1}{\alpha}f^{*'}(\phi_e(\rho)c_e)\phi_e'(\rho), \forall\rho\in[0,1). \tag{11}$$

When $\rho = 0$, we have $\phi_e(0) - \frac{1}{c_e} \geq \frac{1}{\alpha}f^{*'}(\phi_e(0)c_e)\phi_e'(0) = \frac{1}{\alpha}(1 - \frac{1}{\sqrt{\phi_e(0)c_e}})\phi_e'(0)$. After factorization, we have $(\phi_e(0) + \sqrt{\frac{\phi_e(0)}{c_e}} - \frac{\phi_e'(0)}{\alpha})\cdot(\sqrt{\phi_e(0)} - \frac{1}{\sqrt{c_e}}) \geq 0$. A feasible solution is $\phi_e(0) = \frac{1}{c_e}$.

To determine the appropriate right boundary condition, the first step is to identify the maximum utilization possible for any link. Notice that, for any link e, compared to the case where any routing path contains e and other links, it is most likely that e reaches its maximum possible utilization when all agents are routed by a path that only contains e and all agents are of the highest value density $\bar{p}$. An algorithm should ensure that any link's maximum utilization will not be too close to 1 so that the increase of its link cost is less than or equal to the increase of the value collected from agents joined. Define the maximum utilization as the effective utilization $\bar{\rho}_e$. From its definition, we know that $\bar{\rho}_e$ is the maximum ρ that satisfies $\bar{p}dx \geq f'(\rho)\frac{dx}{c_e}$, and $f'(\bar{\rho}_e) = \bar{p}c_e$ follows. Moreover, when $\bar{p} \leq \phi_e(\bar{\rho}_e)$, the algorithm avoids from exceeding the effective utilization, because there is no agent with a value density $\frac{v_i}{r_i} > \phi_e(\bar{\rho}_e) \geq \bar{p}$. Thus, to keep consistent with the effective utilization, we have $\phi_e(\bar{\rho}_e) \geq \bar{p}$, where $\bar{\rho}_e = f'^{-1}(\bar{p}c_e) = f^{*'}(\bar{p}c_e)$. What remains is to construct the worst-case instance and find the conditions on ϕ_e for the worst-case instance to be competitive.

We have observed that any link reaches its effective utilization when it only constitutes unit-length routing paths. Consider the following instance consisting of two phases: In the first phase, for each link e, there come agents with value density gradually increasing from 0 to $\bar{p}$, and they all request link e. In the second phase, for each link e, there come agents with value density of $\bar{p} - \epsilon$. In both phases, the total demand for link e from agents with each value density is larger than the link capacity c_e. The offline optimal solution for this instance only consists of agents in the second phase, and generates a social welfare of $\sum_e[(\bar{p} - \epsilon)\bar{\rho}_e c_e - f(\bar{\rho}_e)] = \sum_e[f^*(\bar{p}c_e) - \epsilon\bar{\rho}_e c_e]$. For $\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$, agents in the first phase will join until the utilization is $\bar{\rho}_e$, and those in the second phase will leave, generating a welfare of $\sum_e \left[\int_0^{\bar{\rho}_e} c_e\phi_e(s)ds - f(\bar{\rho}_e)\right]$. For $\forall \rho \in [0, \bar{\rho}_e]$, taking integral from 0 to ρ at both sides of Eq. (11), we have

$$\begin{aligned}\int_0^{\rho} \phi_e(s)ds - \frac{1}{c_e}(f(\rho) - f(0)) &\geq \frac{1}{\alpha c_e} \cdot (f^*(\phi_e(\rho)c_e) - f^*(\phi_e(0)c_e)) \\ &= \frac{1}{\alpha c_e} \cdot f^*(\phi_e(\rho)c_e). \qquad (12)\end{aligned}$$

When $\rho = \bar{\rho}_e$, it follows from Eq. 12 that

$$\int_0^{\bar{\rho}_e} c_e\phi_e(s)ds - f(\bar{\rho}_e) \geq \frac{1}{\alpha}f^*(\phi_e(\bar{\rho}_e)c_e) > \frac{1}{\alpha}[f^*(\phi_e(\bar{\rho}_e)c_e) - \epsilon\bar{\rho}_e c_e].$$

It then follows from $\phi_e(\bar{\rho}_e) \geq \bar{p}$ that $\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$ is α-competitive for the constructed worst-case instance. □

Theorem 2 reinforces the significance of Eq. (6) by showing that the existence of an α-competitive online mechanism is equivalent to the existence of a solution to the integral version of Eq. (6).

Theorem 2 (Necessity)**.** *For any $\alpha > 0$, if there exists an α-competitive deterministic online mechanism (not necessarily PPMs), then the integral version of Eq.* (6)*, namely, Eq.* (12) *with equality, has at least one solution.*

Proof. The following observation is made: $\forall \rho_e \geq 0, f'(\rho_e) \geq f'(0) = 1$, a unit increase of the utilization level of link e leads to $c \coloneqq f'(\rho_e) \geq 1$ units increase of the cost. Thus, for any online algorithms (containing the online running of the offline optimal algorithm), any agent that produces positive welfare should satisfy $v_i \geq \sum_e \delta_{ij}^e \int_{\rho_e^{(i-1)}}^{\rho_e^{(i-1)}+r_i^e} f'(s)ds \geq \sum_e \delta_{ij}^e r_i^e f'(0) = \sum_e \delta_{ij}^e r_i^e$, where $r_i^e = r_i/c_e$, and without ambiguity, $\rho_e^{(i-1)}$ is the utilization level of any online algorithms before ith arrival in general, and j is the path chosen by the online algorithm for agent i.

The discussion below is limited to deterministic algorithms. Consider the case when the routing path sets of all agents contain only one available routing path for each agent ($|\mathcal{O}_i| = 1, \forall i$), and all routing paths consist of only one link e, $\forall e \in \mathcal{E}$. Denote a group of agents with value density ν and total demand $c_e f^{*'}(\nu c_e)$ as G_ν. Consider the following instance I_p indexed by p, $p \in [0, \bar{p})$: there come G_νs with ν increasing from 0 to p continuously. After that, there comes G_ν with $\nu = p - \epsilon$. The optimal solution is composed of all agents in the last group of the instance I_p, i.e., group $G_{p-\epsilon}$, and its welfare is $(p - \epsilon)c_e f^{*'}((p-\epsilon)c_e) - f(f^{*'}((p-\epsilon)c_e)) = f^*((p-\epsilon)c_e)$. For the instance constructed before, define the utilization of link e of any α-competitive online algorithm after processing G_ν as $\psi_e(\nu)$. The notation of ψ_e is slightly abused here, and the corresponding online algorithm will be clear in the context. The following claim is made: given any α-competitive algorithm, another online algorithm can be found, which is at least α-competitive and incurs a ψ_e function with $\psi_e(\bar{p}) = \bar{\rho}_e$ and $\psi_e(\frac{1}{c_e}) = 0$. Such an algorithm can be found in the following way.

Because $\psi_e(\nu)$ denotes the allocation after processing group G_ν in instance I_p, it is non-negative, non-decreasing in ν, i.e., $\psi_e(\nu) \geq \psi_e(\frac{1}{c_e}) \geq 0$, for all $\nu \in [\frac{1}{c_e}, \bar{p}]$. If $\psi_e(\frac{1}{c_e}) > 0$, it means that agents with $v_i/r_i \leq \frac{1}{c_e}$ will join, however, an online algorithm that discourages those agents from joining will have a better competitive ratio because those agents incur negative welfare ($v_i < \frac{r_i f'(0)}{c_e} = \frac{r_i}{c_e}$) in the system, and thus there always exists a comparably competitive online algorithm with $\psi_e(\frac{1}{c_e}) = 0$. To find an algorithm with $\psi_e(\bar{p}) = \bar{\rho}_e$, we separate the cases when $\psi_e(\bar{p}) > \bar{\rho}_e$ and $\psi_e(\bar{p}) < \bar{\rho}_e$. If $\psi_e(\bar{p}) > \bar{\rho}_e$, we can always construct an algorithm at least α-competitive by stopping the allocation right before the utilization hits the effective utilization $\bar{\rho}_e$, because the increase of the link costs after exceeding the effective utilization is greater than the increase of the value; if $\psi_e(\bar{p}) < \bar{\rho}_e$, we can always allocate the remaining $\bar{\rho}_e - \psi_e(\bar{p})$ of link e to $I_{\bar{p}}$ and achieve a competitive ratio no worse than α. Thus, we find an α-competitive algorithm with $\psi_e(\bar{p}) = \bar{\rho}_e$ and $\psi_e(\frac{1}{c_e}) = 0$. Denote the output of this algorithm as ALG. Given the α-competitiveness, the following inequality holds for $\forall p \in (1/c_e, \bar{p}]$:

$$ALG = \int_{1/c_e}^{p} \nu c_e d\psi_e(\nu) - f(\psi_e(p)) \geq \frac{1}{\alpha} OPT = \frac{1}{\alpha} f^*((p-\epsilon)c_e). \tag{13}$$

For any α-competitive online algorithm, there is a ψ_e that satisfies Eq. (14):

$$\begin{cases} \int_{1/c_e}^{p} \nu d\psi_e(\nu) - \frac{1}{c_e} f(\psi_e(p)) \geq \frac{1}{\alpha c_e} f^*(pc_e), \forall p \in (1/c_e, \bar{p}) \\ \psi_e(\frac{1}{c_e}) = 0, \psi_e(\bar{p}) = \bar{\rho}_e. \end{cases} \tag{14}$$

It is shown in the following that there exists a strictly increasing solution $\underline{\psi_e}$ to Eq. (14) with equality, in other words, we can find an α-competitive online algorithm whose allocation function $\underline{\psi_e}$ for the instance I_p is strictly increasing. Define $\underline{\psi_e}(\nu)$ as the infimum over all feasible solutions to Eq. (14):

$$\underline{\psi_e}(\nu) = \inf \{\psi_e(\nu) | \psi_e \text{ is non-decreasing and feasible for Eq. (14) }\}.$$

The infimum exists because it is evident that a feasible solution to Eq. (14) is bounded from below.

Lemma 3. $\underline{\psi_e}$ *is feasible for Eq.* (14) *with the equality holds and is strictly increasing.*

Proof. By the definition of $\underline{\psi_e}$, it is the greatest lower bound of all feasible ψ_es, and $\underline{\psi_e}(\frac{1}{c_e}) = 0, \underline{\psi_e}(\bar{p}) = \bar{\rho}_e$. Suppose that $\underline{\psi_e}$ is not a feasible solution to Eq. (14), we have $\int_{1/c_e}^{\bar{p}} \nu d\underline{\psi_e}(\nu) - \frac{1}{c_e} f(\underline{\psi_e}(\bar{p})) < \frac{1}{\alpha c_e} f^*(\bar{p}c_e)$. By integration by parts, we have $\nu \underline{\psi_e}(\nu)\Big|_{1/c_e}^{\bar{p}} - \int_{1/c_e}^{\bar{p}} \underline{\psi_e}(\nu) d\nu - \frac{1}{c_e} f(\underline{\psi_e}(\bar{p})) < \frac{1}{\alpha c_e} f^*(\bar{p}c_e)$. We can always push $\underline{\psi_e}$ down on interval $(\frac{1}{c_e}, \bar{p})$ and ensure that it is still non-decreasing until the new function $\tilde{\psi_e}$ is feasible again: $\nu \tilde{\psi_e}(\nu)\Big|_{1/c_e}^{\bar{p}} - \int_{1/c_e}^{\bar{p}} \tilde{\psi_e}(\nu) d\nu - \frac{1}{c_e} f(\tilde{\psi_e}(\bar{p})) \geq \frac{1}{\alpha c_e} f^*(\bar{p}c_e)$. It contradicts that $\underline{\psi_e}$ is the infimum. Thus, $\underline{\psi_e}$ is feasible.

Assume that there exists $p \in (\frac{1}{c_e}, \bar{p})$ such that $\underline{\psi_e}$ satisfies Eq. (14) with a strict inequality. Denote $L(p) = \int_{1/c_e}^{p} \nu d\underline{\psi_e}(\nu) - \frac{1}{c_e} f(\underline{\psi_e}(p))$, $R(p) = \frac{1}{\alpha c_e} f^*(pc_e)$. It means that $L(p) > R(p)$. Let p_0 be the smallest among all such p. If we slightly decrease $\underline{\psi_e}(p_0)$, i.e., $d\underline{\psi_e}(p_0) = -\delta < 0$, then $L(p_0)$ and $R(p_0)$ changes by $\Delta L(p_0) \coloneqq p_0 d\underline{\psi_e}(p_0) - \frac{1}{c_e} f'(\underline{\psi_e}(p_0)) d\underline{\psi_e}(p_0) = (\frac{1}{c_e} f'(\underline{\psi_e}(p_0)) - p_0)\delta, \Delta R(p_0) = 0$, respectively. When δ is very small, it is likely that $L(p_0) + \Delta L(p_0) \geq R(p_0)$ still holds, and a new ψ_e with $\psi_e(p_0) = \underline{\psi_e}(p_0) - \delta$ is also feasible, which contradicts the definition that $\underline{\psi_e}$ is the infimum over all feasible solutions. Thus, $\underline{\psi_e}$ satisfies Eq. (14) with equality.

Assume that there exists $p \in (\frac{1}{c_e}, \bar{p})$ such that $\underline{\psi_e}$ is not strictly increasing on $[p, \bar{p}]$. It means that there exist $p_1, p_2 \in [p, \bar{p}]$ with $p_1 < p_2$ and $\underline{\psi_e}(p_1) = \underline{\psi_e}(p_2)$. In this case, $L(p_1) = L(p_2)$ and $R(p_1) < R(p_2)$, which contradicts that $\underline{\psi_e}$ satisfies Eq. (14) with equality. Thus, $\underline{\psi_e}$ is strictly increasing.

□

Because $\underline{\psi_e}$ is strictly increasing, its inverse function $\underline{\psi_e}^{-1}$ exists. Construct φ_e as follows: for any $p \in (\frac{1}{c_e}, \bar{p})$, $\varphi_e(\rho) = \underline{\psi_e}^{-1}(\rho) = p, \forall \rho \in (0, \bar{\rho}_e)$, $\varphi_e(0) = \frac{1}{c_e}$, $\varphi_e(\bar{\rho}_e) = \bar{p}$. By replacing ν with $\varphi_e(s)$ in Eq. (14), we have $\int_0^{\rho} \varphi_e(s)ds - \frac{1}{c_e} f(\rho) = \frac{1}{\alpha c_e} f^*(\varphi_e(\rho)c_e), \forall \rho \in (0, \bar{\rho}_e)$, which shows that φ_e is a solution to Eq. (12) with equality. □

As a non-autonomous differential equation with singular boundary conditions ($\phi_e(0) = \frac{1}{c_e}$), Eq. (6) is notoriously difficult to analyze. We resort to find the smallest α such that a solution exists numerically, and show its logarithmic growth w.r.t. $\bar{p}$ in Figure 1(a). Here we provide more intuition on the relationship between Eq. (6) and the smallest competitive ratio α. The two boundary conditions require that the two end points (at origin and at the effective utilization) have a minimum function value difference, so that the parameter α in Eq. (6) must be lower bounded because it controls the increasing speed of ϕ. An α too small will lead to an Eq. (6) with no solution. The following corollary characterizes the performance limit of $\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$ when ϕ is analytic.

Corollary 4. *For any analytic ϕ, the best competitive ratio achievable by $\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$ is* 4.

Proof. Assume that there exists a solution to Eq. (6) with α_0. Because $\phi_e(y)$ is analytic, we have

$$\begin{aligned}
\phi_e'(0) &= \lim_{y\to 0+} \phi_e'(y) \\
&= \lim_{y\to 0^+} \alpha_0 \frac{\phi_e - \frac{1}{c_e(1-y)^2}}{1-(c_e\phi_e)^{-1/2}} \\
&\overset{(a)}{=} \lim_{y\to 0^+} \alpha_0 \frac{\phi_e' - \frac{2}{c_e(1-y)^3}}{\frac{1}{2}c_e^{-1/2}\phi_e^{-3/2}\phi_e'} \\
&\overset{(b)}{=} 2\alpha_0 \frac{\phi_e'(0) - 2/c_e}{c_e^{-1/2}\phi_e(0)^{-3/2}\phi_e'(0)} = 2\alpha_0 \frac{\phi_e'(0) - 2/c_e}{c_e\phi_e'(0)}.
\end{aligned}$$

Equality (a) follows from the L'Hôpital's rule, and Equality (b) follows from the continuity of ϕ_e' and ϕ_e at $y = 0$. By multiplying the denominator at both sides, we have $c_e(\phi_e'(0))^2 - 2\alpha_0\phi_e'(0) + \frac{4\alpha_0}{c_e} = 0$. The quadratic equation above should have at least one real solution, leading to $\Delta = 4\alpha_0^2 - 16\alpha_0 \geq 0$, and thus $\alpha_0 \geq 4$. In conclusion, the best competitive ratio achievable by $\mathsf{PPM\text{-}PS}_{\boldsymbol{\phi}}$ is 4.

□

3.1. Analysis of Multiple-The-Index Pricing

It is reported in [11] that the multiple-the-index (MTI) pricing scheme achieves the optimal competitive ratio when the cost function is a power function or the marginal cost function is concave. Specifically, $\phi_e(\rho) = f'(\rho)$ is the optimal price function for $f(\rho) = a\rho^{\gamma+1}$, and $\phi_e(\rho) = f'(2\rho)$ is optimal for cost functions with a concave marginal f'. However, there is no competitive pricing scheme in the existing literature for our case where the cost function has a convex marginal but is not polynomial, and we cannot find an analytical solution to the differential equation in Theorem 1. Thus, it is of interest to study how optimal pricing schemes for other cost functions, i.e., MTI pricing scheme, works in our case. We find that it is very far from the optimal competitive ratio achieved by Algorithm 1. The following theorem shows the performance limit of the MTI pricing scheme.

Theorem 5. *Given $\bar{p}$, the minimum competitive ratio that any MTI pricing scheme can achieve is given by:*

$$\alpha^*_{MTI} = \max_e \frac{2k_e}{k_e^2 - 1} + \frac{2k_e}{k_e + 1} \cdot \frac{(k_e + 2) - (2k_e + 1)\bar{\rho}_e}{(1 - k_e\bar{\rho}_e)(2 - (k_e + 1)\bar{\rho}_e)},$$

where k_e is a solution to the cubic equation $\bar{\rho}_e^2 k^3 + (\bar{\rho}_e^2 - 2\bar{\rho}_e - 2)k + (4 - 2\bar{\rho}_e) = 0$ and $\bar{\rho}_e = 1 - \frac{1}{\sqrt{\bar{p}c_e}}$.

Proof. The k-the-index pricing scheme sets the price at utilization ρ to be proportional to the cost at utilization $k\rho$. Based on Theorem 1, as long as ϕ_e satisfies the BVP in Equation (6), it will lead to a competitive algorithm. Let $\phi_e(0) = \frac{1}{c_e}$, thus $\phi_e(\rho) = \frac{1}{c_e(1-k\rho)^2}$, and it is easy to see that $\phi_e(\bar{\rho}_e) \geq \bar{p}$ when $k \in [1, \frac{1}{\bar{\rho}_e}]$. Given the value density's upper bound $\bar{p}$ and the effective capacities $\bar{\rho}_e = 1 - \frac{1}{\sqrt{\rho c_e}}$, for each link e, there must exists a constant $\alpha_{k,e}$ such that

$$\alpha_{k,e} \geq \frac{\phi_e'(\rho)\left(1 - (c_e\phi_e(\rho))^{-1/2}\right)}{\phi_e(\rho) - \frac{1}{c_e(1-\rho)^2}} = \frac{\frac{2k}{c_e(1-k\rho)^3} \cdot k\rho}{\frac{1}{c_e(1-k\rho)^2} - \frac{1}{c_e(1-\rho)^2}}, \forall \rho \in (0, \bar{\rho}_e),$$

$$\begin{aligned}
\alpha_{k,e} &\geq \max_{\rho \in (0,\bar{\rho}_e)} \frac{\frac{2k}{(1-k\rho)^3} \cdot k\rho}{\frac{1}{(1-k\rho)^2} - \frac{1}{(1-\rho)^2}} \qquad (15)\\
&= \max_{\rho \in (0,\bar{\rho}_e)} \frac{2k^2\rho(1-\rho)^2}{(1-k\rho)(1-\rho)^2 - (1-k\rho)^3}\\
&= \max_{\rho \in (0,\bar{\rho}_e)} \frac{2k^2(1-\rho)^2}{(1-k\rho)[(1-k^2)\rho + 2(k-1)]}\\
&= \max_{\rho \in (0,\bar{\rho}_e)} \frac{2k^2}{k-1} \cdot \frac{(1-\rho)^2}{(1-k\rho)[2-(k+1)\rho]}\\
&= \max_{\rho \in (0,\bar{\rho}_e)} \frac{2k^2}{k-1} \left[\frac{1}{k(k+1)} + \frac{(k-1)[(k+2)-(2k+1)\rho]}{k(k+1)(1-k\rho)[2-(k+1)\rho]}\right]\\
&= \max_{\rho \in (0,\bar{\rho}_e)} \frac{2k}{k^2-1} + \frac{2k}{k+1} \cdot \frac{(k+2)-(2k+1)\rho}{(1-k\rho)(2-(k+1)\rho)}. \qquad (16)
\end{aligned}$$

We can show that $\frac{(k+2)-(2k+1)\rho}{(1-k\rho)(2-(k+1)\rho)}$ is monotonically non-decreasing in $\rho \in (0, \bar{\rho}_e)$ by proving that its first-order derivative is nonnegative, which is omitted together with the proof of $\alpha_{k,e}$'s convexity later. Thus, the competitive ratio of

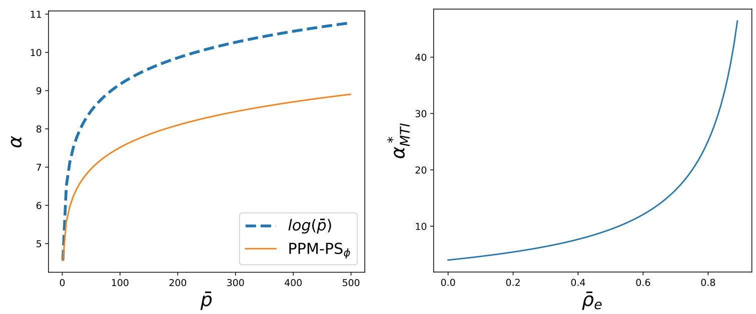

(a) (b)

Figure 1: (a) Best competitive ratio of $\mathsf{PPM\text{-}PS}_\phi$ vs. $\bar{p}$. Link capacity is set to 40. (b) Best competitive ratio of the MTI pricing vs. $\bar{\rho}_e$.

the k-the-index pricing scheme is

$$\alpha_k = \max_e \alpha_{k,e} = \max_e \frac{2k}{k^2-1} + \frac{2k}{k+1} \cdot \frac{(k+2)-(2k+1)\bar{\rho}_e}{(1-k\bar{\rho}_e)(2-(k+1)\bar{\rho}_e)}. \tag{17}$$

It can be shown that $\alpha_{k,e}$ is convex in k by proving that the second-order derivative of $\alpha_{k,e}$ with respect to k is non-negative. Setting the first-order derivative of $\alpha_{k,e}$ with respect to k to zero, we find that $\alpha_{k,e}$ reaches its minima when k is the second largest real solution to the following cubic equation

$$\bar{\rho}_e^2 k^3 + (\bar{\rho}_e^2 - 2\bar{\rho}_e - 2)k + (4 - 2\bar{\rho}_e) = 0. \tag{18}$$

□

Figure 1(b) shows that when capacity is 40 and $\bar{p}$ is 5, α^*_{MTI} is over 40 ($\bar{\rho}_e \approx 0.93$) while $\mathsf{PPM\text{-}PS}_\phi$ is around 5-competitive (see Figure 1(a)). This indicates that when the cost function is generally convex but not polynomial, the MTI pricing is much worse than our proposed online mechanism $\mathsf{PPM\text{-}PS}_\phi$.

4. Conclusions and Future Works

In this work, we formulated the online path-aware path selection problem as a mechanism design problem and focused on developing competitive posted-pricing mechanisms. We established sufficient and necessary conditions on the pricing scheme to ensure competitiveness and discussed the fundamental performance limits of posted price mechanisms. We also analyzed the MTI pricing scheme in terms of its best competitive ratio and showed that it is numerically inferior to our proposed scheme. It could be beneficial to employ our modeling and analysis approach to enable multipath routing. Additionally, conducting beyond-worst-case analysis for the addressed problem is also worth exploring.

Acknowledgement

This work is supported in part by the Guangzhou Municipal Science and Technology Project under Grant 2023A03J0011, Guangdong Provincial Key Laboratory of Integrated Communications, Sensing and Computation for Ubiquitous Internet of Things, and National Foreign Expert Project, Project Number G2022030026L.

References

[1] D. Barrera, L. Chuat, A. Perrig, R. M. Reischuk, and P. Szalachowski, "The SCION internet architecture," *Communications of the ACM*, vol. 60, no. 6, pp. 56–65, 2017.

[2] S. Scherrer, M. Legner, A. Perrig, and S. Schmid, "Incentivizing stable path selection in future internet architectures," *ACM SIGMETRICS Performance Evaluation Review*, vol. 48, no. 3, pp. 12–13, 2021.

[3] S. Dawkins, "Path aware networking: Obstacles to deployment (a bestiary of roads not taken)," in *Internet Engineering Task Force, Internet-Draft draft-irtf-panrg-what-not-to-do-00*, 2020.

[4] Y.-J. Chen and J. Zhang, "Design of price mechanisms for network resource allocation via price of anarchy," *Mathematical Programming*, vol. 131, no. 1-2, pp. 333–364, 2012.

[5] X. Tan, B. Sun, A. Leon-Garcia, Y. Wu, and D. H. K. Tsang, "Mechanism design for online resource allocation: A unified approach," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 4, no. 2, pp. 1–46, 2020.

[6] L. Yang, A. Zeynali, M. H. Hajiesmaili, R. K. Sitaraman, and D. Towsley, "Competitive algorithms for online multidimensional knapsack problems," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 5, no. 3, pp. 1–30, 2021.

[7] X. Tan, A. Leon-Garcia, Y. Wu, and D. H. K. Tsang, "Online combinatorial auctions for resource allocation with supply costs and capacity limits," *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 4, pp. 655–668, 2020.

[8] B. Awerbuch, Y. Azar, and S. Plotkin, "Throughput-competitive on-line routing," in *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pp. 32–40, IEEE, 1993.

[9] A. Blum, A. Gupta, Y. Mansour, and A. Sharma, "Welfare and profit maximization with production costs," in *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pp. 77–86, 2011.

[10] N. Buchbinder, J. S. Naor, *et al.*, "The design of competitive online algorithms via a primal–dual approach," *Foundations and Trends® in Theoretical Computer Science*, vol. 3, no. 2–3, pp. 93–263, 2009.

[11] Z. Huang and A. Kim, "Welfare maximization with production costs: A primal dual approach," *Games and Economic Behavior*, vol. 118, pp. 648–667, 2019.