Fair opportunistic schedulers for Lossy Polling systems

Vartika Singh
IEOR, IIT Bombay, India, vsvartika@iitb.ac.in

Veeraruna Kavitha
IEOR, IIT Bombay, India, vkavitha@iitb.ac.in

***Abstract*—Polling systems with losses are useful mathematical objects that can model many practical systems like travelling salesman problem with recurrent requests. One of the less studied yet an important aspect in such systems is the disparity in the utilities derived by the individual stations. Further, the random fluctuations of the travel conditions can have significant impact on the performance. This calls for a scheduler that caters to the fairness aspect, depends upon the travel conditions and the dynamic system state.**

Inspired by the generalized alpha-fair schedulers of wireless networks, we propose a family of schedulers that further considers binary knowledge of the travel conditions. These schedulers are opportunistic, allocate the server to a station with bad travel condition only when the station has accumulated too little a utility by the decision epoch. We illustrate that the disparities among the individual utilities diminish to zero, as fairness factor increases, and further that the price of fairness decreases as the number of stations increase.

***Index Terms*—opportunistic, alpha-fair schedulers, price of fairness, losses.**

I. Introduction

The motivation for this study comes from a variant of travelling salesman problem, where the salesman keeps travelling from one location to another to serve the waiting customers. An important aspect is an optimal (preferably) dynamic policy that guides the salesman about the next location to be visited after each service. It is natural that some locations will be preferred more by the salesman in comparison to others, possibly because of travel conditions, demands, etc. This may lead to the starvation of less preferred locations and calls for a fair policy.

We consider a single server that visits several stations to serve the demands. After each service, it makes a decision to either remain in the same station, or to travel to other stations depending upon travel conditions, system state, etc. In all, we have a polling system (e.g., [1], [8], [10]) with controlled (dynamic) service and random switching times.

Fairness is a relatively less studied aspect in polling systems (e.g., [10]) and our focus is on utilities of individual stations and further on opportunistic aspects. We take inspiration from generalised α-fair opportunistic scheduler of the wireless networks (e.g., [5], [7]) which considers fair allocation of resources (like bandwidth) among users over time; authors in [3] argue that the price of fairness (PoF) of opportunistic schedulers for wireless networks with large users is negligible. We propose a Fair opportunistic Polling Scheduler (FoPS) that prescribes the next station to be visited by the server, depending upon current server location, travel conditions and number of customers waiting at different stations; to cater to fairness aspects, FoPS also depends upon the accumulated utilities of the individual stations till the decision epoch. The FoPS maintains a balance between the efficiency and the fairness dictated by a single parameter α. Basically, it chooses (non-empty) stations with good travel conditions as long as none of the stations are starved.

The first author's work is partially supported by PMRF, India.

The salient features and contributions of our work are: i) we propose a consolidated objective function that considers losses (occurred by losing the customers) and the rewards (obtained by serving customers); ii) the proposed scheduler is based on some anticipated utilities of the individual stations that depend upon the current state and travel conditions; iii) we define a measure of fairness (MoF) that quantifies the disparity in utilities of all the stations and show that the metric reduces with increase in α or the number of stations; and finally, iv) the PoF, i.e., the price of fairness (a metric similar to price of anarchy) defined in [3], [9], is miraculously negligible even with heterogeneous demands.

To summarize, the proposed schedulers achieve various levels of fairness (the most fair policy results in almost equal individual station-utilities), with almost negligible loss of efficiency in several scenarios. In other scenarios, the price of fairness reduces to a value that can be achieved in an ideal system, where the travel conditions are always good and this is possible with large number of stations. Surprisingly, the efficient scheduler among FoPS is itself fair with large number of stations. We perform numerical simulations to reaffirm our findings, and observe around ten stations suffice for negligible MoF, this number further decreases with increased α. Interestingly, PoF is negligible in all the case studies.

Polling systems are well studied objects in literature (see [1] for an elaborate and recent survey) and we discuss a few relevant strands here; authors in [10] study fairness aspects with respect to various service disciplines (within queues) like exhaustive, gated, etc.; [8] considers dynamic scheduling policies to schedule jobs within each queue. However, to the best of our knowledge, none of these papers consider fairness aspects across stations and opportunistic advantage resulting with some minimal knowledge about switching conditions. Further, they do not discuss the price and measure of fairness.

II. Problem statement and background

There are m stations and a single server. At any station i, the demands arrive according to an independent Poisson process with rate λ_i and join the queue, Q_i. The server travels from one station to another and keeps serving the waiting customers. The service times are independent and identically distributed (i.i.d.) across all the stations with mean $1/\mu_s$. Further the buffer length at station i is b_i, i.e., if a new demand arrives at station i, and if the queue length is b_i, such a demand is lost.

We consider a controlled service, in which the server upon arrival to station i serves exactly one customer (if the queue is non-empty); it then makes a decision either to stay in station i or to travel to some other station, depending upon on the system state. If the queue Q_i is empty, the server waits for some time and then makes a new decision; for mathematical simplicity, we model the waiting time distribution to be the same as the service time distribution. The server does not serve any demand during wait times. The state of the system is primarily described by the number of customers waiting at different stations and the travel conditions. From any station to any other station, the travel conditions can either be good or bad. When the travel conditions are good (respectively bad), the travel time is distributed according to a known distribution (same for all stations) with mean μ_g (respectively μ_b); naturally $\mu_g < \mu_b$. The travel condition from station i to j is good with probability $0 < p_{ij} < 1$.

In all, we have a polling system (e.g., [1]) with controlled (dynamic) service and random switching times. Further, the controlled dynamic decision also depends upon the travel conditions to the rest of the stations; relevant travel conditions are assumed to be known at the decision epochs. In our model, depending upon the controlled decision, the server can be idle and wait for next arrival in the queue even if there are demands at other stations. Furthermore, we consider fairness aspect, that attempts to reduce the disparities in individual station utilities.

Fairness is a well studied concept in the context of wireless networks (see [3], [4] etc.). Our approach is to adopt the idea from wireless networks to design a fair policy. Before we discuss these details, we digress briefly to discuss the fairness concepts in wireless networks.

A. Motivation: Wireless α-fair scheduler

In a wireless network, a scheduler allocates the resource to one of the users in each time-slot/epoch, and aims to maximize the sum of the time-averages ($\{\bar{u}_i\}$) of the resultant utilities of all the users. When time-slot k is allocated to user i, it achieves instantaneous utility $U_{i,k}$, zero else. The achievable utilities $\{U_{i,k}\}_k$ of any user i are i.i.d. and different users can have different distributions. The scheduler can observe the utilities before allocating the resource. To maximize[1] the sum of time-average utilities, an efficient scheduler allocates the resource to the one with the highest utility in each epoch. This leads to starvation of some of the users over time, especially ones with lower expected utilities. In such scenarios, it becomes important to be fair while deriving the optimal resource allocation policy $\boldsymbol{\beta}$.

Generalized α-fair schedulers maximize a concave function of time-average utilities (by Ergodic theorem, expected utilities) of all the users, and achieve various levels of fairness depending upon parameter α (e.g., [5], [7]):

$$\sup_{\boldsymbol{\beta}=(\beta_1,\cdots,\beta_m)} \sum_i \Gamma_\alpha\left(\bar{u}_i(\boldsymbol{\beta})\right) \text{ with } \bar{u}_i(\boldsymbol{\beta}) := E[U_i\beta_i(\mathbf{U})]],$$

$$\Gamma_\alpha(\bar{u}) := \frac{\bar{u}^{1-\alpha}\mathbb{1}_{\{\alpha\neq 1\}}}{1-\alpha} + \log(\bar{u})\mathbb{1}_{\{\alpha=1\}}, \qquad (1)$$

where $\mathbf{U} := (U_1, \cdots U_m)$. In the above, max-min fairness is achieved with $\alpha \to \infty$, proportional fairness is achieved with $\alpha = 1$, and $\alpha = 0$ achieves efficient scheduler.

The solution $\boldsymbol{\beta}^{(\alpha)} = (\beta_1, \ldots, \beta_m)$ of (1), for any α, is given by the solution $\bar{\mathbf{u}}^{(\alpha)} = (\bar{u}_1, \ldots, \bar{u}_m)$ of following m-dimensional fixed point equation ([5], [7]):

$$\bar{u}_n = E[U_n\beta_n(\mathbf{U})], \ \beta_n(\mathbf{U}) := \prod_{i\neq n} \mathbb{1}\left\{\frac{U_n}{(\bar{u}_n)^\alpha} \geq \frac{U_i}{(\bar{u}_i)^\alpha}\right\}. \qquad (2)$$

An iterative algorithm that uses the average utilities derived by the individual users till epoch k, represented by $\bar{\mathbf{U}}_k = (\bar{U}_{1,k}, \cdots, \bar{U}_{m,k})$, is proposed to attain the above solution:

$$\begin{aligned} \bar{U}_{n,k+1} &= \bar{U}_{n,k} + \frac{1}{k+1}\left(U_{n,k+1}\beta_{n,k+1} - \bar{U}_{n,k}\right), \\ \beta_{n,k+1} &= \prod_{i\neq n} \mathbb{1}\left\{\frac{U_{n,k+1}}{(\bar{U}_{n,k})^\alpha} \geq \frac{U_{i,k+1}}{(\bar{U}_{i,k})^\alpha}\right\}. \end{aligned} \qquad (3)$$

In [5], [7], it is proved that the above algorithm converges weakly to asymptotic utilities (2) for i.i.d. channels. Thus the data scheduler $\boldsymbol{\beta}_k$, for slot k, depends on the channel utilities $\mathbf{U}_k = (U_{1,k}, \ldots, U_{m,k})$ as well as the average utilities $\bar{\mathbf{U}}_{k-1}$.

Further, there must be a price of fairness for deviating from the efficient scheduler. But the results in [3] show that the price of fairness is negligible if there are a large number of users. We would also investigate this aspect.

III. Fair Dynamic Opportunistic schedulers

The idea is to derive a fair policy for polling systems following the guidelines from wireless networks. In our case, server is the resource to be allocated to the users (stations). We need to define the instantaneous utilities of the stations in such a way that, the utility of the server is the sum of instantaneous utilities of all the stations over time. The instantaneous utility from station i should consist of gains from serving the customers at i and losses due to customer drops because of buffer length b_i. We begin with some definitions.

Let $\mathbf{N}_k = (N_{1,k}, \cdots, N_{m,k})$ be the number of customers waiting at various stations, and $S_{k-1} \in \{1, \cdots, m\}$ be the server location at the beginning of k-th decision epoch (i.e., immediately after $(k-1)$-th service/waiting). The flag $C_{i,k} \in \{g, b\}$ represents the travel condition, if the server has to travel

[1] Recall the utilities are independent across the epochs and are identical for the same user. Thus efficient scheduler maximizes sum of the expected utilities of all the users, which by Ergodic theorem, almost surely equals the sum of the time averages. As we will see soon, in our context the utilities are not independent across epochs; thus to optimize the sum of the expected utilities in strict sense, one requires Markov Decision Process (MDP) framework. This would be an interesting future aspect that we would investigate.

from S_{k-1} to i. Set $C_{i,k} = 0$ if $i = S_{k-1}$, to indicate no travel is required, and let $\mathbf{C}_k = (C_{1,k}, \cdots, C_{m,k})$.

The scheduler in (3) depends upon the average utilities $\bar{\mathbf{U}}_k$ derived by individuals till the decision epoch k. Observe from (3) that the users with smaller $\bar{U}_{i,k}$ have higher weight factor; higher α magnifies to a larger extent. *The scheduler also considers opportunistic aspect as the scheduling decision additionally depends upon the instantaneous utilities* $\mathbf{U}_k = \{U_{i,k}\}_i$. Our aim is similar, average utilities as defined below can again be used for similar purpose, opportunities provided by the travel conditions $\mathbf{C}_k$ and the current state, $(S_{k-1}, \mathbf{N}_k)$ can also be used: first define the average utility obtained till the decision epoch k for any Q_i (factor $\gamma > 1$ helps in mathematical tractability),

$$\bar{U}_{i,k} = \bar{U}_{i,k-1} + \frac{1}{k^\gamma}\left(\hat{U}_{i,k} - \bar{U}_{i,k-1}\right), \tag{4}$$

where $\hat{U}_{i,k}$ has to be appropriately defined to reflect the opportunities provided by $\mathbf{Z}_k := (\mathbf{X}_k, \mathbf{C}_k)$ with $\mathbf{X}_k = (S_{k-1}, \mathbf{N}_k, \bar{\mathbf{U}}_{k-1})$; we will soon see that $\hat{U}_{i,k}$ will be anticipated utilities of some $U_{i,k}$, the actual instantaneous utilities gained by the decision taken at epoch k; we begin with defining $U_{i,k}$ while $\hat{U}_{i,k}$ are defined in (6).

As already mentioned, the utility of any station is comprised of gains and losses and depends upon the scheduling decision. Say $\boldsymbol{\beta}_k = \boldsymbol{\beta}(\mathbf{Z}_k) = (\beta_{1,k}, \ldots, \beta_{m,k})$ is the decision, where $\beta_{j,k}$ is the probability of travelling from $Q_{S_{k-1}}$ to Q_j, with $\beta_{j,k} \in [0,1]$ and $\sum_j \beta_{j,k} = 1$. *Station i gains a reward w if one of it's users is served.* It's losses are captured by the number of users that have overflown during the travel time (if any) and the service time. Thus the instantaneous utility of station i because of decision $\boldsymbol{\beta}_k$ equals,

$$U_{i,k}(\boldsymbol{\beta}_k, \mathbf{Z}_k) = w\mathbb{1}_{\{S_k=i\}}\mathbb{1}_{\{\tilde{N}_{i,k}\ge 1\}} - \sum_j \mathbb{1}_{\{S_k=j\}} L_i(S_{k-1}, j), \tag{5}$$

where $\tilde{N}_{i,k}$ is the number of customers waiting in Q_i just before the service starts (recall $N_{i,k}$ is the number immediately after the previous service and exactly equals $\tilde{N}_{i,k}$ if the server does not travel), L_i represents the losses in queue Q_i during the time interval between two decision epochs and the probability $P(S_k = j|\boldsymbol{\beta}_k) = \beta_{j,k}$ for any j. The losses L_i depend upon the service times, travel conditions (travel times) and hence upon S_{k-1} (previous server location) and the next station S_k chosen by scheduler $\boldsymbol{\beta}_k$.

We immediately have the following observation. The state $\{\mathbf{Z}_k\}$ (in fact $\{\mathbf{X}_k\}$) evolves as a non-homogeneous Markov Chain with countable state space under any given dynamic policy $\boldsymbol{\beta}$ that depends only upon the system state $\mathbf{Z}$.

A. Opportunistic utilities based on anticipation

It is clear that the utilities in (5) are not known at the decision epoch k. Hence we propose to use the expected values of $\{U_{i,k}\}_i$, conditioned on the available information $\mathbf{Z}_k$. These form instantaneous utilities for k-th decision epoch and equal,

$$\hat{U}_{i,k}(\boldsymbol{\beta}, \mathbf{Z}_k) = \beta_{i,k}\hat{\mathbb{I}}_i(\mathbf{Z}_k) - \sum_j \beta_{j,k}\hat{l}_i(j, \mathbf{Z}_k), \tag{6}$$

where i) $\hat{\mathbb{I}}_i(\mathbf{Z}_k) := wE[\tilde{N}_{i,k} > 0|\mathbf{Z}_k, S_k = i]$ equals the expected gain at Q_i by serving a customer; and ii) $\hat{l}_i(j, \mathbf{Z}_k) := E[L_i(S_{k-1}, S_k)|\mathbf{Z}_k, S_k = j]$ is the expected loss at Q_i (due to buffer length b_i), when station j is chosen for next service.

By conditioning on $\mathbf{Z}_k = \mathbf{z} = (\mathbf{x}, \mathbf{c})$, the expected losses and gains of station i are given by,

$$\begin{aligned} \hat{\mathbb{I}}_i(\mathbf{z}) &= w\left(\mathbb{1}_{\{n_i>0\}} + \mathbb{1}_{\{n_i=0, s\neq i\}} E\left[\mathbb{1}_{\{\mathcal{N}_i(T_{s,i})>0\}}|c_i\right]\right), \\ \hat{l}_i(j, \mathbf{z}) &= E\left[(\mathcal{N}_i(T_{s,j}+J) + n_i - b_i)^+ |c_j, n_i\right], \end{aligned} \tag{7}$$

where $\mathcal{N}_i(T_{s,i})$ is the number of Poisson arrivals in Q_i during the travel time $T_{s,i}$ and J is the service time (waiting time if queue is empty). Further the anticipated losses of any station are bounded uniformly by $\hat{l}^* := \max_q\{\lambda_q(\mu_b + 1/\mu_s)\}$, as:

$$\begin{aligned} \hat{l}_i(j, \mathbf{z}) &= E[(\mathcal{N}_i(T_{s,j}+J) + n_i - b_i)^+ |c_j, n_i], \\ &\le E[\mathcal{N}_i(T_{s,j}+J)|c_j] \le \lambda_i(\mu_b + 1/\mu_s) \le \hat{l}^*. \end{aligned} \tag{8}$$

B. Fair opportunistic Polling Scheduler (FoPS(α))

We are now ready to describe the proposed scheduler. Taking inspiration from wireless scheduler given in (2) and (3), we propose the following scheduler, parameterized by α, for decision epoch k when state $\mathbf{Z}_k = \mathbf{z} = (\mathbf{x}, \mathbf{c})$,

$$\textbf{FoPS}(\alpha): \qquad \boldsymbol{\beta}_k^*(\mathbf{z}) = \arg\max_{\boldsymbol{\beta}} \sum_i \frac{\hat{U}_{i,k}(\boldsymbol{\beta}, \mathbf{z})}{\tilde{\bar{u}}_i^\alpha}, \tag{9}$$

where $\tilde{\bar{u}}_i := \max\{\delta, \bar{u}_i\}$; here $\delta > 0$ is a small quantity introduced to take care of possible negative values. Observe that this scheduler is lot more complicated than the scheduler in (3): a) in wireless networks, if a user is not chosen at an epoch, it's instantaneous utility is zero, however in our system the unselected stations get negative utility; b) further the negative utility depends upon travel condition of the chosen path (station); c) thus the utility of any station in any decision epoch just does not depend upon whether the station is selected or not; d) and hence the system evolves according to a non-homogeneous Markov chain driven by dynamic decision (9).

We can re-write the decision (9) as,

$$\begin{aligned} \boldsymbol{\beta}_k(\mathbf{z}) &= \arg\max_{\tilde{\boldsymbol{\beta}}} \sum_i \tilde{\beta}_i O_i(\mathbf{z}) \quad \text{where,} \\ O_i(\mathbf{z}) &:= \frac{\hat{\mathbb{I}}_i(\mathbf{z})}{\tilde{\bar{u}}_i^\alpha} - \sum_j \frac{\hat{l}_j(i, \mathbf{z})}{\tilde{\bar{u}}_j^\alpha}. \end{aligned} \tag{10}$$

In the above expression, O_i can be interpreted as the *fair-weighted anticipated utility of the server* upon choosing Q_i. Hence, the decision $\boldsymbol{\beta}_k(\mathbf{z})$ is given by,

$$\beta_{q,k}(\mathbf{z}) = \frac{\mathbb{1}_{\{q \in \arg\max_i O_i(\mathbf{z})\}}}{|\arg\max_i O_i(\mathbf{z})|}, \text{ for any } q. \tag{11}$$

We conclude this section with a conjecture that describes the limits of the time-averages in (4), which is required for Theorem 2. Further the conjecture provides interesting and important insights which could be of independent interest. The conjecture is verified using several numerical examples (e.g., Figure 1) and is partially supported by Theorem 1 which also proves the convergence of utilities of (4). Towards this we require some important definitions.

Let $\mathbf{Y}_k(\bar{\mathbf{u}}) := (S_{k-1}, \mathbf{N}_k, \mathbf{C}_k)$ be the chain whose transitions are similar to the transitions of corresponding components of $\mathbf{Z}_k$ except that $\boldsymbol{\beta}_k$ in (11) is now defined using fixed vector $\bar{\mathbf{u}}$. *Observe that* $\mathbb{Y}(\bar{\mathbf{u}}) = \{\mathbf{Y}_k(\bar{\mathbf{u}})\}_k$ *is a finite-state Markov chain, can be reducible or irreducible depending upon* $\bar{\mathbf{u}}$ and will have unique stationary distribution in each irreducible closed class for any $\bar{\mathbf{u}}$. Further, it satisfies the standard Ergodic theorem (e.g., [6]) for any $\bar{\mathbf{u}}$. These Ergodic chains, in our opinion, form the basis to derive the almost sure (a.s.) limit of time-average utilities defined in (4) under an assumption that uses the following definition.

D.0 The vector $\bar{\mathbf{u}}^*$ is said to be Ergodic-interior if there exists an $\epsilon > 0$, such that $\mathbb{Y}(\bar{\mathbf{u}})$ has transition probabilities same as $\mathbb{Y}(\bar{\mathbf{u}}^*)$ for all $\bar{\mathbf{u}} \in \{\bar{\mathbf{u}}' : |\bar{\mathbf{u}}' - \bar{\mathbf{u}}^*| < \epsilon\}$.

Theorem 1: **[Ergodicity]** Assume $\gamma > 1$. For any α and i, the time-average utilities,

$$\bar{V}_{i,k} := \frac{1}{k}\sum_{t=1}^{k} U_{i,t}, \tag{12}$$

and the utilities defined in (4) converge. The respective limits $\bar{\mathbf{v}}, \bar{\mathbf{u}}$ can depend upon sample path ω, and further:

$$\bar{v}_i \;=\; E\left[\hat{U}_i(\boldsymbol{\beta}, \mathbf{Z})\right] \text{ for any } i, \tag{13}$$

where the expectation is under[2] a stationary distribution (SD) of Markov chain $\mathbb{Y}(\bar{\mathbf{u}})$, if $\bar{\mathbf{u}}$ is Ergodic-interior as defined in **D**.0. The transitions of Markov chain $\mathbb{Y}(\bar{\mathbf{u}})$ are captured by:

$$\beta_q(\mathbf{Y}) \;=\; \frac{\mathbb{1}_{\{q \in \arg\max_i O_i(\mathbf{Y},\bar{\mathbf{u}})\}}}{|\arg\max_i O_i(\mathbf{Y},\bar{\mathbf{u}})|}, \text{ with} \tag{14}$$

$$O_i(\mathbf{Y},\bar{\mathbf{u}}) \;:=\; \frac{\hat{\mathbb{I}}_i(\mathbf{Y},\bar{\mathbf{u}})}{\tilde{\bar{u}}_i^{\alpha}} - \sum_j \frac{\hat{l}_j(i,\mathbf{Y},\bar{\mathbf{u}})}{\tilde{\bar{u}}_j^{\alpha}}, \; \tilde{\bar{u}}_i := \max\{\delta, \bar{u}_i\},$$

$$\hat{U}_i(\boldsymbol{\beta},\mathbf{Y},\bar{\mathbf{u}}) \;=\; \beta_i \hat{\mathbb{I}}_i(\mathbf{Y},\bar{\mathbf{u}}) - \sum_j \beta_j \hat{l}_i(j,\mathbf{Y},\bar{\mathbf{u}}).$$

Proof is in Appendix. ■

A conjecture: When γ is close to 1, we anticipate the two limits of the above theorem $\bar{\mathbf{u}}$ and $\bar{\mathbf{v}}$ to be close (observe $\bar{\mathbf{v}}$ also equals the time limit of (4) but with $\gamma = 1$, as $\hat{\mathbf{U}}$ are conditional expectations of $\mathbf{U}$). For now we conjecture them to be approximately equal, further using (13) conjecture that the following fixed point equation is satisfied as $\gamma \downarrow 1$,

$$\bar{u}_i \;=\; E[\hat{U}_i(\boldsymbol{\beta}, \mathbf{Z})] \text{ for any } i. \tag{15}$$

Remarks: i) The proof of the conjecture is immediate for $\alpha = 0$ by the well-known Ergodic theorem provided in standard text books (see also [6]) when $\gamma = 1$; $\{\mathbf{Z}_k\}$ is then finite state space irreducible Markov chain from (10) under **A.1** (given in section IV). ii) With $\alpha > 0$, $\{\mathbf{Z}_k\}$ is a non-homogeneous Markov chain and we are working towards the proof of the conjecture for this case. For this short paper, we continue with further analysis by assuming the conjecture to be true. iii) If the conjecture were true, the average utilities of (4) converge to the limit utilities $\bar{\mathbf{u}} = \bar{\mathbf{u}}^{(\alpha)}$, which are given by stationary expected values of an appropriate Markov chain $\mathbb{Y}$ with 'limit transition probabilities' given in (14); these are governed by fixed point equation (15) as the limit transition probabilities depend upon the limit utilities $\bar{\mathbf{u}}$. *iv) the conjecture is used only in the proof of Theorem 2, while the rest of the results are derived independently.*

[2]For Markov chains with more than one irreducible class, the chain eventually gets absorbed into one of them depending upon the sample path, and the average (12) converges to the corresponding stationary expected value.

IV. Fairness Analysis

Our aim in this section is to derive the analysis of the proposed FoPS(α) schedulers. We begin with defining two measures that respectively capture the 'amount of fairness' achieved and the 'price paid for that fairness'.

A scheduler is said to be efficient if it maximizes the server utility $\sum_i \bar{u}_i$. From (9), FoPS(0) is efficient among the FoPS(α) schedulers. The server utility is maximized by efficient scheduler, but the individual utilities of the stations $\{\bar{u}_i\}$ can have *huge disparities depending upon the travel condition statistics and the arrival rates of individual stations. The precise aim of the fair schedulers is to reduce this disparity.* However, there is an anticipated reduction in the server utility and hence efficiency. The price of fairness (PoF) captures this price (as in [3], [9]):

$$\mathbb{P}(\alpha) := \frac{\sum_{j=1}^m \bar{u}_j(0) - \sum_{j=1}^m \bar{u}_j(\alpha)}{\sum_{j=1}^m \bar{u}_j(0)}, \tag{16}$$

where $\{\bar{u}_j(\alpha)\}$ is solution of (15) (or the limit of (4)) for given α. One can say a scheduler is max-min fair (see [5] and references therein) if it ensures that the utilities of all the stations are equal. Hence we define the measure of fairness (MoF) of any scheduler using the amount of disparities between the individual utilities as below:

$$\mathbb{M}(\alpha) = \max_{i,j} \frac{\tilde{\bar{u}}_i(\alpha)}{\tilde{\bar{u}}_j(\alpha)} - 1. \tag{17}$$

By definition $\mathbb{M} \geq 0$, the smaller the value, the 'fairer' is the system. Further, $\mathbb{M}$ can also be zero if all the utilities $\bar{u}_i < \delta$. We make the following assumption for further analysis,

A.1 Maximum possible losses should be less than maximum possible gain, i.e.,

$$\rho_B := \max_{i,q} \sum_j \mathcal{N}_j E[(T_{i,q} + J)\,|c_q = b] < w.$$

The first observation about the FoPSs is that all the stations are served *infinitely often (i.o.)*, i.e., the number of visits to any station increases to ∞ (proof is in Appendix):

Lemma 1: Define the event, $\mathcal{I} := \cap_i \{S_t = i \text{ i.o.}\}$, where $\{S_t = i \text{ i.o.}\} = \cap_t \cup_{k \geq t} \{S_k = i\}$ and assume **A**.1. Then for any $\alpha > 0$, we have $P(\mathcal{I}) = 1$. ■

Thus, with $\alpha > 0$, all the stations are served infinitely many times, however the more important aspect is the measure of fairness. Towards this, we have the next result, which shows that the measure of fairness reduces, i.e., 'fairness' increases as α increases. The proof (in Appendix) of this theorem uses the Conjecture 1, and hence is true subjected to the verification of the latter. Nonetheless, the simulations in the next section illustrate that $\mathbb{M}$ indeed decreases with an increase in α.

Theorem 2: **[MoF decreases with** α**]** There exists a $1 \leq B < \infty$ such that $\mathbb{M} \leq B^{1/\alpha} - 1$ for every $\alpha > 0$. ■

From (11), FoPS(α) with $\alpha > 0$ would not allocate server to a bad-station (i.e., one with bad travel condition) unless the latter has accumulated too little a utility by the decision epoch and hence are also opportunistic in nature; they attempt to maintain a balance between fairness and efficiency. From Theorem 2,

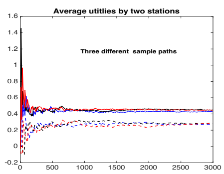

Fig. 1. All three sample paths converge to the same limit

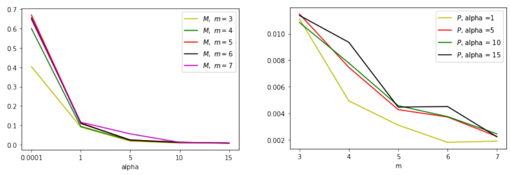

Fig. 2. **MoF $\mathbb{M}$ and PoF $\mathbb{P}$ versus α and m:** $\lambda = 0.67/m$, $\mu_g = 2$, $\mu_b = 6$, $\sigma_g = \sigma_b = 0.1$, $b = 5$, $\mu_s = 1/3$, $p_{ij} \in \{0.1, 0.9\}$, $w = 6$.

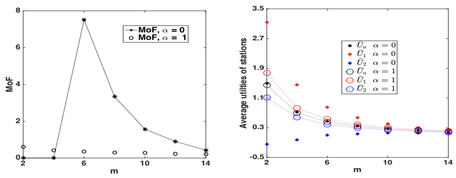

Fig. 3. **Good stations with more demands:** $\mu_g = 2$, $\sigma_g = 0.01$, $\mu_b = 6$ $\sigma_b = .08$, $b_i \equiv 4$, $\mu_s = 1/3$ $p_{i,j} \in \{0.1, 0.9\}, \lambda_i \in \{0.37/m, 0.93/m\}$ and $w = 4$

as $\alpha \to \infty$, the measure of fairness $\mathbb{M}$ decreases to zero, implying almost identical utilities for all the stations; observe this is true even for the case with heterogeneous demands. However one now has to understand PoF, and this is possible if probably the MoF of the efficient scheduler is analyzed.

Fairness of Efficient scheduler: An efficient scheduler maximizes system utility and is known to be unfair in the context of communication networks (e.g., [3], [5]). The systems appear to follow a different logic, when the schedulers are opportunistic; we have an interesting observation about the fairness of efficient scheduler itself: authors in [3] argue that the PoF (16) for opportunistic wireless schedulers decreases as m increases and our polling systems results mirror the same.

The efficient scheduler usually attempts to schedule stations with good travel condition, and does not pay attention to the stations that have accumulated low utility (and have been starved). However for large m, we show that there always exist good non-empty stations.

Lemma 2: Let $\mathcal{G}_k := \{i : N_{i,k} > 0, C_{i,k} \in \{g, 0\}\}$. Then for any k, $P(\text{support}(\beta_k^\alpha) \subset \mathcal{G}_k) \to 1$ as $m \to \infty$, when $\alpha = 0$.

Proof: At any decision epoch consider, if possible, any $i \in \mathcal{G}$ and any $j \in \mathcal{G}^c$. Then,

$$O_i(s) - O_j(s) = w + \sum_{k=1}^{m} \left(\hat{l}_k(j, s) - \hat{l}_k(i, s) \right) - \hat{\mathbb{I}}_j(s), \quad > w - \hat{\mathbb{I}}_j(s) \geq 0,$$

as either $\hat{\mathbb{I}}_j(s) = 0$ or $C_j = b$. Thus any $j \notin \mathcal{G}$ is not scheduled, if the set $\mathcal{G}$ is non-empty. The probability, $P(|\mathcal{G}_k| = 0) \geq P\left(\cap_{i \leq m} \{C_i = b, \mathcal{N}_i(J) = 0\}\right) \to 0$ as $m \to \infty$ for any k. ■

In view of the above Lemma, one can anticipate that the MoF/PoF of the system with large m is close to that of an ideal system with only good travel conditions. If further $\lambda_i = \lambda$ for all i, then $\mathbb{M}$ of the ideal system is zero – from equation (15), $\bar{u}_i$ is the same for all i by symmetry for any given α. Thus *with large m, the MoF of efficient scheduler itself is small.* Further the ideal system has identical station utilities for all α, hence one can anticipate that the PoF is also zero.

V. Numerical results

We now present few examples to illustrate the variations in measures $\mathbb{M}$ and $\mathbb{P}$ for different levels of fairness, number of stations, and with possibly heterogeneous demands. In the first example (Figure 2), travel towards two stations (from any starting point) is in bad condition with probability 0.1, and for the remaining with probability 0.9; travel time follows Normal distribution (parameters in caption). MoF $\mathbb{M}$ is high for small α, and decreases with increase in α, for any m as discussed in Theorem 2. Further, $\mathbb{P}$ (right sub-figure) is smaller for higher α, more importantly it is negligible for all the cases.

In another example in Figure 3, half the stations are good stations, i.e., have good travel conditions with high probability, and half are bad, further, the bad stations have lesser demand (λ). In the right sub-figure, we plot the average utility of good and bad stations as a function of m, for different α. When there are only two stations, with $\alpha = 0$, the bad station has a negative utility; this goes in line with the fact that efficient scheduler prefers only good stations. With $\alpha > 0$, both the stations have positive utility and disparity is lesser. As m increases, the utility of a bad station approaches that of a good station even for $\alpha = 0$ (blue and red curves with only markers). *This reaffirms the fairness of efficient scheduler in a system with large m.* In the left sub-figure, we observe that $\mathbb{M}$ is smaller with higher α as in Theorem 2; further, it also decreases with m. Interestingly, server utility (average of individual station utilities) is almost the same for $\alpha = 0$ and $\alpha = 1$; hence from (16) the PoF is negligible. However MoF $\mathbb{M}$ is negligible only for $m > 10$. Thus the PoF/MoF is negligible even for the case with heterogeneous demands either with large α or with large number of stations.

Acknowledgement. The Masters thesis of Mr Rishabh Kumar has inspired us to write this paper.

VI. Conclusion

We consider polling systems with losses (due to customers lost), gains (by serving) and opportunities (scheduler has some information about travel conditions). Assuming just binary knowledge of the travel conditions, we propose a family of fair schedulers that can achieve any desired level of fairness. Fairness is measured in terms of disparities in utilities of the individual stations. Using partial theoretical arguments and a conjecture, we illustrate the following: i) when the fairness parameter α increases, the proposed scheduler achieves max-min fairness, i.e., the disparity in individual utilities reduces to zero; ii) as the number of stations increases, the efficient scheduler itself achieves the max-min fairness and the price of fairness reduces. The numerical examples illustrate that sometimes the schedulers are fair without paying much price, for number of stations as less as ten.

The work is incomplete without the proof of the conjecture. Further, the dynamic decisions in queuing systems should also consider the future cost as in Markov Decision Process framework.

References

[1] Borst, Sem, and Onno Boxma. "Polling: past, present, and perspective." Top 26.3 (2018): 335-369.
[2] H. Kushner, Harold, and G. Yin. Stochastic approximation and recursive algorithms and applications. Springer Science & Business Media, 2003.
[3] Veeraruna Kavitha, Nandyala Hemachandra, and Mayur Zambre. "Opportunistic schedulers and asymptotic price for fairness." Computer Communications 148 (2019): 27-41.
[4] Kavitha, Veeraruna, Nandyala Hemachandra, and Debayan Das. "Fairness via priority scheduling." 2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE, 2013.
[5] Veeraruna Kavitha, Altman, E., El-Azouzi, R. and Sundaresan, R., 2013. Fair scheduling in cellular systems in the presence of noncooperative mobiles. IEEE/ACM Transactions on Networking, 22(2), pp.580-594.
[6] Vassiliou, P.C., 2020. Laws of Large numbers for non-homogeneous Markov systems. Methodology and Computing in Applied Probability.
[7] Kushner, H.J. and Whiting, P.A., 2004. Convergence of proportional-fair sharing algorithms under general conditions. IEEE transactions on wireless communications, 3(4), pp.1250-1259.
[8] Wierman, Adam, Erik MM Winands, and Onno J. Boxma. "Scheduling in polling systems." Performance Evaluation 64.9-12 (2007): 1009-1028.
[9] Bertsimas, Dimitris, Vivek F. Farias, and Nikolaos Trichakis. "The price of fairness." Operations research 59.1 (2011): 17-31.
[10] Shapira, Gil, and Hanoch Levy. "On fairness in polling systems." Annals of Operations Research (2016): 1-33.

Appendix

Proof of Theorem 1: Step 1. Using law of large numbers and the fact that conditional utilities are bounded in either direction, i.e. (see (8)),

$$-\hat{l}^* \leq \hat{U}_{i,k} \leq w \text{ for all } i, k \text{ and sample paths } \omega,$$

one can prove that the time averages $\bar{U}_{i,k}$ in (4) are uniformly bounded between $[-\hat{l}^*, w]$, for all $k \geq K(\omega)$ and for almost all sample paths ω, with $\hat{l}^*$ as in (8).

Step 2. From (4) and (6), again with $\Delta := \max\{\hat{l}^*, w\} + w$,

$$|\bar{U}_{i,k+m} - \bar{U}_{i,k}| \leq \sum_{q=1}^{m} \frac{\Delta}{(k+q)^\gamma} \text{ for all } k \geq K(\omega),$$

Thus $\{\bar{U}_{i,k}\}$ is a Cauchy sequence (a.s.) and hence converges to a limit $\bar{u}_i \in [-l^*, w]$ a.s., for every i. Observe that the limit $\bar{\mathbf{u}} = (\bar{u}_1, \cdots, \bar{u}_m)$ depends upon sample path ω.

Step 3. From (11), the transition probabilities of the chain $\mathbf{Y}_k := (S_{k-1}, \mathbf{N}_k, \mathbf{C}_k)$ freeze for all $k > K(\omega)$ (redefine $K(\omega)$ if required to ensure this) as $\bar{\mathbf{U}}_k \to \bar{\mathbf{u}}$; this is because the finitely many decisions in (11) remain the same for all $\bar{\mathbf{U}}_k \in \{\bar{\mathbf{u}}' : |\bar{\mathbf{u}}' - \bar{\mathbf{u}}| \leq \epsilon\}$ for some appropriate $\epsilon > 0$. This step is true under **D**.0.

Step 4. Thus $\{\mathbf{Y}_k\}$ is a homogeneous finite-state Markov chain for all $k \geq K(\omega)$. We will have at least one irreducible closed class and a stationary distribution (SD). Observe $\bar{\mathbf{v}}_k$ of (12) equals time-average of an appropriate function of $\mathbf{Y}_k$; hence it converges to the corresponding stationary expectation (depends upon $\mathbf{Y}_k$ at the time transition probabilities freeze) by standard Ergodic theorems applicable for finite-state Markov chains. ∎

Proof of Lemma 1: Proof is by contradiction. If say under FoPS the server stops visiting station i after some n epochs, i.e., say $S_k \neq i$ for all $k \geq n$. Then from (4) and (5) there exist a finite τ, such that $\bar{\mathbf{U}}_{i,t-1} < \delta$ and $N_{i,t} \geq 1$ for all $t \geq \tau$ with δ as in (9); so $\hat{\mathbb{I}}_i(\mathbf{Z}_t) = w$ for all $t \geq \tau$.

Consider a $t \geq \tau$, for which $N_{S_{t-1},t} = 0$ (the queue at current server location is empty), $C_{i,t} = g$; such a t always exists because of the non-zero probability associated with all the relevant events. Consider $O_i(\mathbf{Z}_t) - O_j(\mathbf{Z}_t)$ where j is any station that is visited i.o., obviously one such station always exists, (dropping notation $\mathbf{Z}_t$, see (10) and say $S_{t-1} = s$):

$$O_i(s) - O_j(s) = \frac{w}{\delta^\alpha} - \frac{\hat{\mathbb{I}}_j(s)}{\tilde{\bar{u}}_j^\alpha} + \sum_l \frac{\hat{l}_l(j,s) - \hat{l}_l(i,s)}{\tilde{\bar{u}}_l^\alpha}.$$

When $j = s$ from (7) we have $\hat{\mathbb{I}}_j(s) = 0$ and then (as $\tilde{\bar{u}}_j \geq \delta$)

$$O_i(s) - O_j(s) \geq \frac{w - \sum_{l=1}^m \hat{l}_l(i,s)}{\delta^\alpha} > 0, \qquad (18)$$

by assumption **A**.1. For $j \neq s$, $\hat{l}_l(j,s) \geq \hat{l}_l(i,s)$ since $C_{i,t} = g$,

$$O_i(s) - O_j(s) \geq \frac{w}{\delta^\alpha} - \frac{\hat{\mathbb{I}}_j(s)}{\tilde{\bar{u}}_j^\alpha} \geq 0,$$

and then since the ties are broken randomly and equally likely, station i would be chosen at some such t from (11), thus a contradiction. ∎

Proof of Theorem 2: We provide the proof by contradiction. For the given α, say $(\bar{u}_1, \ldots, \bar{u}_m)$ is the fixed point of (15). Without loss of generality, assume that $\bar{u}_1 <, \ldots, < \bar{u}_m$ and if possible, say $\frac{\bar{u}_m}{\bar{u}_1} > B^{1/\alpha}$ where $B > 1$ is such that,

$$B > \frac{w}{\min_{i,j,k} E[\mathcal{N}_k(T_{i,j}) | C_j = g]}.$$

We drop $\mathbf{Z}$ from the notation for simplicity and retain only S and say, $S = s$ with $s \neq m$. Consider $O_s(s) - O_m(s)$,

$$\begin{aligned}
&O_s(s) - O_m(s) \\
&= \frac{\hat{\mathbb{I}}_s(s) - \hat{l}_s(s,s) + \hat{l}_s(m,s)}{\tilde{\bar{u}}_s^\alpha} - \frac{\hat{\mathbb{I}}_m(s) - \hat{l}_m(m,s) + \hat{l}_m(s,s)}{\tilde{\bar{u}}_m^\alpha} \\
&\quad + \sum_{j \neq s,m} \frac{\hat{l}_j(m,s) - \hat{l}_j(s,s)}{\tilde{\bar{u}}_j^\alpha}, \\
&= \frac{1}{\tilde{\bar{u}}_m^\alpha} \left(\frac{\tilde{\bar{u}}_m^\alpha}{\tilde{\bar{u}}_s^\alpha} \hat{\mathbb{I}}_s(s) - \hat{\mathbb{I}}_m(s) + \sum_{j=1}^m \frac{\tilde{\bar{u}}_m^\alpha}{\tilde{\bar{u}}_j^\alpha} (\hat{l}_j(m,s) - \hat{l}_j(s,s)) \right), \\
&\geq \frac{1}{\tilde{\bar{u}}_m^\alpha} \left(\frac{\tilde{\bar{u}}_m^\alpha}{\tilde{\bar{u}}_1^\alpha} (\hat{l}_1(m,s) - \hat{l}_1(s,s)) - \hat{\mathbb{I}}_m(s) \right) > 0,
\end{aligned}$$

by the choice of B. Hence, $\beta_m(\mathbf{Z}) = 0$ if $S \neq m$. This implies $P(S' = m | S \neq m) = 0$, which in turn implies that the state space is reducible under limit transition probabilities, with class of states $\{S \neq m\}$ being closed. If $P(S = m) = 0$, then from (15), $\bar{u}_m < 0$ which contradicts $\tilde{\bar{u}}_m > \tilde{\bar{u}}_1 B \geq \delta B$. Thus $P(S = m) > 0$ (this is under SD), which implies $\{S = m\}$ is also closed class under limit transition probabilities, and further that $P(S = m) = 1$; this also implies $P(S \neq m) = 0$. Hence under the SD, $N_i = b_i$ and $\tilde{\bar{u}}_i = \delta$ for all $i < m$. Now the limiting chain only schedules m-th station, and thus by **A**.1 (and standard arguments for queues with load factor less than one) $P(N_m = 0) > 0$. But as argued in (18), with $N_1 > 0$, $N_m = 0$ and $\tilde{\bar{u}}_1 = \delta$ we have $O_1 \geq O_m$, which in turn implies $P(S = 1) > 0$. This is again a contradiction and hence the assumption that $\mathbb{M} > B^{1/\alpha} - 1$ is not true. ∎