This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Aalto, Samuli; Lassila, Pasi

Whittle index approach to energy-aware dispatching

Published in:
International Teletraffic Congress

DOI:
[10.1109/ITC30.2018.00012](https://doi.org/10.1109/ITC30.2018.00012)

Published: 01/01/2018

Document Version
Peer reviewed version

Please cite the original version:
Aalto, S., & Lassila, P. (2018). Whittle index approach to energy-aware dispatching. In *International Teletraffic Congress* (pp. 19 - 27). IEEE. <https://doi.org/10.1109/ITC30.2018.00012>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Whittle index approach to energy-aware dispatching

Samuli Aalto

Department of Communications and Networking
Aalto University, Finland
samuli.aalto@aalto.fi

Pasi Lassila

Department of Communications and Networking
Aalto University, Finland
pasi.lassila@aalto.fi

Abstract—A data center can be modeled as a set of parallel queues, and the dispatcher decides to which queue the arriving jobs are routed. We consider an energy-aware dispatching system in a Markovian setting, where each server upon becoming empty enters a sleep mode to save energy and to activate the server after sleep incurs an additional setup delay cost. We seek to optimize the performance-energy trade-off by applying the so-called Whittle index approach. As our main result, we rigorously prove, under a certain technical assumption, that the problem is indexable, and derive the explicit form of the Whittle index. Our numerical examples demonstrate that the resulting energy-aware Whittle index policy is able to perform very close to the numerically obtained optimal policy.

I. INTRODUCTION

We consider a system of parallel servers with their own queues where an arriving job is dispatched to one of the servers upon arrival. Dispatching (a.k.a. task assignment) problems belong to the optimal control problems for parallel server systems [1]. However, there are very few exact optimality results available. Maybe the most well-known is the optimality of the Join-Shortest-Queue (JSQ) policy in a homogeneous setting where all the servers have the same service rate, originally proved for exponential service times in [2] and thereafter generalized to any service time distributions with a non-decreasing hazard rate in [3].

One approach presented in the literature to produce near-optimal dispatching policies is to utilize the policy iteration algorithm for Markov decision processes. In particular with Poisson arrivals, the first policy iteration (FPI) step can sometimes be made explicitly, see, e.g., [4], [5], [6], [7], [8], [9].

In all papers mentioned above, the servers are assumed to be ordinary ones swapping between busy and idle states. In this paper, we are, however, more interested in *energy-aware servers* that can be switched off (without any delay) to save energy but, on the other hand, incur a setup delay when switched back on. In [10], such servers are called InstantOff servers, while the ordinary ones are referred to as NeverOff servers. In [11], [12], [13], [14], it is shown, under various assumptions, that the optimal sleep state control policy for a single server is either InstantOff or NeverOff. The FPI approach has also been applied for the dispatching problem in the context of such energy-aware InstantOff servers, see, e.g., [15], [16], [17], [18], [19].

In this paper, instead of the FPI approach, we apply the *Whittle index approach* to the energy-aware dispatching problem in order to generate alternative near-optimal control

policies. This approach was originally developed in the context of restless bandits [20], and it has successfully been applied, e.g., to the opportunistic scheduling problems in [21], [22], [23], [24]. Niño-Mora [25] managed to apply the Whittle index approach to the dispatching problem in a successful way when the queues behave as one-dimensional birth-death processes with a finite state space. Argon et al. [6] assumed that the queues behave as ordinary M/M/1 queues with an infinite state space. They also managed to show the indexability property for a large class of cost functions (including the linear holding costs) and derive the corresponding index values. Larrañaga et al. [26] considered dispatching problems where the queues behave as one-dimensional birth-death processes with an infinite state space. They characterized the index assuming that the optimal policy is of threshold type.

Our target in this paper is to derive near-optimal policies for dispatching problems where the queues behave as energy-aware M/M/1 queues provided with InstantOff servers, which is an essentially more complicated task than for the ordinary M/M/1 queues (or more general birth-death processes) due to their two-dimensional state space. To study the performance-energy trade-off in such systems, we assume energy costs in addition to normal linear holding costs. We give a sufficient condition for the system parameters under which the problem is indexable and also determine the corresponding index values explicitly. In addition, our numerical examples demonstrate that the resulting energy-aware Whittle index policy is able to perform very close to the numerically obtained optimal policy.

The energy-aware dispatching problem and the Whittle index approach to tackle it by utilizing a relaxation of the original problem, are described in more detail below in Section II. In Section III, we prove that, at least under a certain condition, the relaxed version is indexable and derive the corresponding Whittle index for an energy-aware M/M/1 system with an InstantOff server. In Section IV, we introduce the energy-aware Whittle index policy, and compare it numerically with the FPI dispatching policy and the ordinary JSQ rule in Section V. Finally, Section VI concludes the paper.

II. ENERGY-AWARE DISPATCHING PROBLEM

We consider the following energy-aware dispatching problem. New jobs arrive according to a Poisson process with rate λ . At the arrival time, the job is dispatched to one of K parallel servers, each provided with an infinite buffer. Each server i is an exponential server with rate μ_i , i.e., the

service time of any job in this server is independently and exponentially distributed with mean $1/\mu_i$. The server is said to be *busy* when it is processing jobs. When server i has processed all the jobs and its buffer becomes empty, it is immediately switched *off*. Server i remains switched-off until a new job is dispatched to it, after which it still needs an exponential *setup* phase with mean $1/\delta_i$, before becoming busy again. In line with [10], such servers are called *InstantOff* servers in this paper.

The state of server i at time t is described by the pair $(N_i(t), Z_i(t))$, where $N_i(t) \in \mathcal{N} = \{0, 1, \dots\}$ denotes the number of customers and $Z_i(t) \in \mathcal{Z} = \{\text{off}, \text{setup}, \text{busy}\}$ the energy state. Let $P_i(z) \geq 0$ denote the (constant) power consumption in energy state z . It is natural to assume that

$$0 \leq P_i(\text{off}) < P_i(\text{setup}) \leq P_i(\text{busy}). \quad (1)$$

In addition, we introduce the following differential notation for $z \in \{\text{setup}, \text{busy}\}$:

$$\hat{P}_i(z) = P_i(z) - P_i(\text{off}) > 0. \quad (2)$$

With each server i and time t , we also associate a decision variable $A_i(t) \in \mathcal{A} = \{0, 1\}$. If $A_i(t) = 1$, then the next arriving customer is dispatched to server i ; otherwise not. Naturally, we require that, for any t ,

$$\sum_{i=1}^K A_i(t) = 1. \quad (3)$$

At time t , server i incurs costs at rate

$$C_i(t) = h_i N_i(t) + \beta P_i(Z_i(t)), \quad (4)$$

where $h_i > 0$ is the holding cost rate per job and $\beta \geq 0$ is an energy weight factor. The problem is to choose the decision variables $A_i(t)$ in such a way that the expected long-run average cost,

$$\lim_{T \rightarrow \infty} E \left[\frac{1}{T} \int_0^T \left(\sum_{i=1}^K (h_i N_i(t) + \beta P_i(Z_i(t))) \right) dt \right], \quad (5)$$

is minimized, subject to constraint (3) for all t . The problem is considered in the context of continuous-time Markov decision processes (CTMDP) [27], where the decision variables $A_i(t)$ can only be changed when the state of the whole system,

$$((N_i(t), Z_i(t)) \mid i \in \{1, \dots, K\}),$$

changes.

Note that, without constraint (3), the problem is trivially solved by choosing passivity: $A_i(t) = 0$ for all i and t . Including constraint (3), however, makes the problem extremely hard.

Following the ideas originally developed by Whittle in [20], we relax the dispatching problem (5) by replacing the strict constraint (3) with an averaged one,

$$\lim_{T \rightarrow \infty} E \left[\frac{1}{T} \int_0^T \left(\sum_{i=1}^K A_i(t) \right) dt \right] = 1, \quad (6)$$

and approach the relaxed problem by the Lagrangian methods. As a result, we get the following *separate* subproblems. For each server i , we try to minimize the objective function

$$\lim_{T \rightarrow \infty} E \left[\frac{1}{T} \int_0^T (h_i N_i(t) + \beta P_i(Z_i(t)) + \nu(1 - A_i(t))) dt \right], \quad (7)$$

where the Lagrangian multiplier ν can be interpreted as the price of passivity per time unit.

We further note that, if $\nu \leq 0$, then there is no need to be active so that $A_i(t) = 0$ is optimal for any t . Thus, from this on, we assume that $\nu \geq 0$.

III. WHITTLE INDEX FOR INSTANTOFF SERVERS

In this section, we consider the optimal solution for the separate subproblems with objective function (7), which is then used to generate a heuristic policy for the original problem (5) subject to the strict constraint (3) in Section IV.

We start by defining the continuous-time Markov decision process related to the separate subproblems with objective function (7). The state of the “system” (i.e., server i) is given by the pair $x = (n, z)$, where $n \in \mathcal{N}$ refers to the number of customers and $z \in \mathcal{Z}$ to the energy state, and the possible actions $a \in \mathcal{A}$ are “to dispatch” ($a = 1$) and “not to dispatch” ($a = 0$). Thus, the state space

$$\mathcal{X} = \{(0, \text{off})\} \cup \{(n, \text{setup}) \mid n \geq 1\} \cup \{(n, \text{busy}) \mid n \geq 1\}$$

is discrete and the action space $\mathcal{A}$ finite. Let $c_i(x, a)$ denote the cost rate in state $x \in \mathcal{X}$ after action $a \in \mathcal{A}$. In our model, we have, for any $x = (n, z) \in \mathcal{X}$ and $a \in \mathcal{A}$,

$$c_i(x, a; \nu) = h_i n + \beta P_i(z) + \nu(1 - a).$$

In addition, let $q_i(y|x, a) \geq 0$ denote the transition intensity from state $x \in \mathcal{X}$ to another state $y \in \mathcal{X} \setminus \{x\}$ after action $a \in \mathcal{A}$. In our model, the following transitions are possible:

$$\begin{aligned} q_i((1, \text{setup})|(0, \text{off}), a) &= a\lambda, \\ q_i((n, \text{busy})|(n, \text{setup}), a) &= \delta_i, & n \geq 1, \\ q_i((n+1, \text{setup})|(n, \text{setup}), a) &= a\lambda, & n \geq 1, \\ q_i((0, \text{off})|(1, \text{busy}), a) &= \mu_i, \\ q_i((n-1, \text{busy})|(n, \text{busy}), a) &= \mu_i, & n \geq 2, \\ q_i((n+1, \text{busy})|(n, \text{busy}), a) &= a\lambda, & n \geq 1. \end{aligned}$$

Since the state space $\mathcal{X}$ is discrete, the action space $\mathcal{A}$ finite, and the cost rate linear with respect to n , there is a stationary deterministic policy π_i^* (described by a function from the state space $\mathcal{X}$ to the action space $\mathcal{A}$) that minimizes the expected average costs (7) [27]. The optimal policy π_i^* is characterized by the *optimality equations* defined for each state $x \in \mathcal{X}$ by

$$\bar{c}_i(\nu) = \min_{a \in \mathcal{A}} \left\{ c_i(x, a; \nu) + \sum_{y \neq x} q_i(y|x, a)(v_i(y; \nu) - v_i(x; \nu)) \right\}, \quad (8)$$

where $\bar{c}_i(\nu)$ denotes the minimum expected average cost rate (per time unit) and $v_i(x; \nu)$ refers to the value function, which

gives the difference in the expected total costs when the optimal stationary policy is applied and the system started from state x or in equilibrium.

In our model, the optimality equations (8) read as follows. For state $x = (0, \text{off})$,

$$\bar{c}_i(\nu) = \beta P_i(\text{off}) + \min\{\nu, \lambda \Delta_i(0, \text{off}; \nu)\}, \quad (9)$$

where we have defined

$$\Delta_i(0, \text{off}; \nu) = v_i(1, \text{setup}; \nu) - v_i(0, \text{off}; \nu).$$

For states $x \in \{(n, \text{setup}) \mid n \geq 1\}$,

$$\begin{aligned} \bar{c}_i(\nu) = & h_i n + \beta P_i(\text{setup}) + \\ & \delta_i(v_i(n, \text{busy}; \nu) - v_i(n, \text{setup}; \nu)) + \\ & \min\{\nu, \lambda \Delta_i(n, \text{setup}; \nu)\}, \end{aligned} \quad (10)$$

where we have defined

$$\Delta_i(n, \text{setup}; \nu) = v_i(n+1, \text{setup}; \nu) - v_i(n, \text{setup}; \nu).$$

For state $x = (1, \text{busy})$,

$$\begin{aligned} \bar{c}_i(\nu) = & h_i + \beta P_i(\text{busy}) + \\ & \mu_i(v_i(0, \text{off}; \nu) - v_i(1, \text{busy}; \nu)) + \\ & \min\{\nu, \lambda \Delta_i(1, \text{busy}; \nu)\}, \end{aligned} \quad (11)$$

where we have defined

$$\Delta_i(1, \text{busy}; \nu) = v_i(2, \text{busy}; \nu) - v_i(1, \text{busy}; \nu).$$

For states $x \in \{(n, \text{busy}) \mid n \geq 2\}$,

$$\begin{aligned} \bar{c}_i(\nu) = & h_i n + \beta P_i(\text{busy}) - \mu_i \Delta_i(n-1, \text{busy}; \nu) + \\ & \min\{\nu, \lambda \Delta_i(n, \text{busy}; \nu)\}, \end{aligned} \quad (12)$$

where we have defined

$$\Delta_i(n, \text{busy}; \nu) = v_i(n+1, \text{busy}; \nu) - v_i(n, \text{busy}; \nu).$$

Thus, by the optimality equations (9)–(12), we deduce that, for any state $x \in \mathcal{X}$, it is optimal to dispatch ($a = 1$) the next job to server i in state x if and only if

$$\nu \geq \lambda_i \Delta_i(x; \nu). \quad (13)$$

We say that the optimization problem with objective function (7) is *indexable*¹ if, for any state $x \in \mathcal{X}$, there exists $\nu_i^*(x) \in [-\infty, \infty]$ such that

- (i) it is optimal to dispatch ($a = 1$) the next job to server i in state x if $\nu \geq \nu_i^*(x)$;
- (ii) it is optimal not to dispatch ($a = 0$) the next job to server i in state x if $\nu < \nu_i^*(x)$.

Such a value $\nu_i^*(x)$ is referred to as the *Whittle index* of state x for the problem with objective function (7).

Below we show that, under assumption

$$\delta_i > \mu_i \left(1 + \frac{\beta}{h_i} \hat{P}_i(\text{setup}) \right), \quad (14)$$

¹Note that we have adapted Whittle's notation of indexability [20] to our dispatching problem in a similar way as done in [6]. As a result, the heuristic index policy for the original problem dispatches the arriving job to the server with the *lowest* index.

the optimization problem with objective function (7) is indexable and derive an explicit expression for the corresponding Whittle index. For that, we introduce the following sets of states, for any $n \geq 1$,

$$\begin{aligned} \mathcal{T}^{0, \text{off}} &= \{(0, \text{off})\}, \\ \mathcal{T}^{n, \text{busy}} &= \{(m, z) \in \mathcal{X} \mid m \leq n-1\} \cup \{(n, \text{busy})\}, \\ \mathcal{T}^{n, \text{setup}} &= \{(m, z) \in \mathcal{X} \mid m \leq n\}. \end{aligned}$$

Note that the family of these sets of states is monotonous,

$$\mathcal{T}^{0, \text{off}} \subset \mathcal{T}^{1, \text{busy}} \subset \mathcal{T}^{1, \text{setup}} \subset \mathcal{T}^{2, \text{busy}} \subset \mathcal{T}^{2, \text{setup}} \subset \dots,$$

defining the following *total order* among all states $x \in \mathcal{X}$:

$$(0, \text{off}), (1, \text{busy}), (1, \text{setup}), (2, \text{busy}), (2, \text{setup}), \dots \quad (15)$$

Let us now define a family of *threshold policies* corresponding to these *activity sets*,

$$\begin{aligned} \pi_i^{0, \text{off}}(x) &= \begin{cases} 1, & \text{if } x \in \mathcal{T}^{0, \text{off}}, \\ 0, & \text{otherwise;} \end{cases} \\ \pi_i^{n, \text{busy}}(x) &= \begin{cases} 1, & \text{if } x \in \mathcal{T}^{n, \text{busy}}, \\ 0, & \text{otherwise;} \end{cases} \\ \pi_i^{n, \text{setup}}(x) &= \begin{cases} 1, & \text{if } x \in \mathcal{T}^{n, \text{setup}}, \\ 0, & \text{otherwise;} \end{cases} \end{aligned} \quad (16)$$

together with a rudimentary threshold policy,

$$\pi_i^\emptyset(x) = 0, \quad x \in \mathcal{X}. \quad (17)$$

The key idea is to prove that, under assumption (14), the optimal policies for the optimization problem with objective function (7) belong to these threshold policies

$$\Pi_i^T = \{\pi_i^\emptyset, \pi_i^{0, \text{off}}\} \cup \{\pi_i^{n, \text{busy}}, \pi_i^{n, \text{setup}} \mid n \geq 1\}. \quad (18)$$

Note that each of these policies generate an irreducible finite-state Markov process, the steady-state distribution of which is rather easily derived based on the global balance equations and the normalization condition.

Before the main result (given below in Theorem 1), we give some auxiliary results related to these threshold policies. The first one (Proposition 1) gives an expression for the corresponding expected long-run average cost rate defined in (7) and denoted by $\bar{c}_i^\pi(\nu)$ for the threshold policies π . The remaining ones (Propositions 2–5) characterize the value function difference $\Delta^\pi(x; \nu)$ for the threshold policies π . Due to lack of space, we have omitted their proofs, which are tedious but straightforward based on the so-called Howard equations for these policies.

Proposition 1: For all threshold policies $\pi \in \Pi_i^T$,

$$\bar{c}_i^\pi(\nu) = h_i \frac{D_i^\pi}{G_i^\pi} + \beta \frac{E_i^\pi}{G_i^\pi} + \nu \left(1 - \frac{F_i^\pi}{G_i^\pi} \right),$$

where D_i^π , E_i^π , F_i^π , and G_i^π are given in Appendix.

Proof: Let $\pi \in \Pi_i^T$ and denote by $P\{X_i^\pi = x\}$ the steady-state distribution of the corresponding irreducible finite-state Markov process. The interpretation of the constants above is as follows: G_i^π is a normalization constant, and

$$\begin{aligned} \frac{D_i^\pi}{G_i^\pi} &= \sum_{n \geq 1} n(P\{X_i^\pi = (n, \text{busy})\} + \\ &\quad P\{X_i^\pi = (n, \text{setup})\}), \\ \frac{E_i^\pi}{G_i^\pi} &= P_i(\text{off})P\{X_i^\pi = (0, \text{off})\} + \\ &\quad P_i(\text{setup}) \sum_{n \geq 1} P\{X_i^\pi = (n, \text{setup})\} + \\ &\quad P_i(\text{busy}) \sum_{n \geq 1} P\{X_i^\pi = (n, \text{busy})\}, \\ \frac{F_i^\pi}{G_i^\pi} &= \sum_{x \in \mathcal{T}^\pi} P\{X_i^\pi = x\}. \end{aligned}$$

The formulas given in Appendix are based on an elementary steady-state analysis of finite-state Markov processes, the details of which are, thus, omitted here. ■

Proposition 2: Assume (14), and denote $\pi = \pi_i^\emptyset$. Then, for any $m \geq 0$,

$$\begin{aligned} \Delta_i^\pi(m+1, \text{busy}; \nu) - \Delta_i^\pi(0, \text{off}; \nu) \\ &= h_i \left(\frac{m+1}{\mu_i} - \frac{1}{\delta_i} \right) - \beta \hat{P}_i(\text{setup}) \frac{1}{\delta_i} > 0, \\ \Delta_i^\pi(m+1, \text{setup}; \nu) - \Delta_i^\pi(0, \text{off}; \nu) \\ &= h_i \frac{m+1}{\mu_i} - \beta \hat{P}_i(\text{setup}) \frac{1}{\delta_i} > 0. \end{aligned}$$

Proposition 3: Denote $\pi = \pi_i^{0, \text{off}}$. Then, for any $m \geq 0$,

$$\begin{aligned} \Delta_i^\pi(m+1, \text{setup}; \nu) - \Delta_i^\pi(1, \text{busy}; \nu) \\ &= h_i \left(\frac{m}{\mu_i} + \frac{1}{\delta_i} \right) > 0, \\ \Delta_i^\pi(m+2, \text{busy}; \nu) - \Delta_i^\pi(1, \text{busy}; \nu) \\ &= h_i \frac{m+1}{\mu_i} > 0. \end{aligned}$$

Proposition 4: Assume (14), and denote $\pi = \pi_i^{n, \text{busy}}$. Then, for any $m \geq 0$,

$$\begin{aligned} \Delta_i^\pi(n+m+1, \text{busy}; \nu) - \Delta_i^\pi(n, \text{setup}; \nu) \\ &= h_i \left(\frac{m+1}{\mu_i} - \frac{1}{\delta_i} \right) > 0, \\ \Delta_i^\pi(n+m+1, \text{setup}; \nu) - \Delta_i^\pi(n, \text{setup}; \nu) \\ &= h_i \frac{m+1}{\mu_i} > 0. \end{aligned}$$

Proposition 5: Denote $\pi = \pi_i^{n, \text{setup}}$. Then, for any $m \geq 0$,

$$\begin{aligned} \Delta_i^\pi(n+m+1, \text{setup}; \nu) - \Delta_i^\pi(n+1, \text{busy}; \nu) \\ &= h_i \left(\frac{m}{\mu_i} + \frac{1}{\delta_i} \right) > 0, \\ \Delta_i^\pi(n+m+2, \text{busy}; \nu) - \Delta_i^\pi(n+1, \text{busy}; \nu) \\ &= h_i \frac{m+1}{\mu_i} > 0. \end{aligned}$$

Note that (14) is needed for Propositions 2 and 4.

Theorem 1: Under assumption (14), the optimization problem with objective function (7) for an InstantOff server i is indexable, and the corresponding Whittle index is given by

$$\nu_i^*(x) = h_i H_i(x) + \beta B_i(x), \quad (19)$$

where we have used the following notations. For $x = (0, \text{off})$,

$$\begin{aligned} H_i(0, \text{off}) &= \frac{\lambda}{\delta_i} + \frac{\lambda}{\mu_i}, \\ B_i(0, \text{off}) &= \hat{P}_i(\text{setup}) \frac{\lambda}{\delta_i} + \hat{P}_i(\text{busy}) \frac{\lambda}{\mu_i}. \end{aligned}$$

For any $x = (n, \text{busy})$, where $n \geq 1$,

$$\begin{aligned} H_i(n, \text{busy}) &= \begin{cases} \frac{b}{b(1+a)-a} \left(\frac{b((n+1)-(n+2)b+b^{n+2})}{(b-1)^2} + \right. \\ \quad \left. a^2 - (n+1)a + a(b-a)\left(\frac{a}{1+a}\right)^n \right), & \lambda \neq \mu_i, \\ \frac{1}{2} (n+1)(n+2) + \\ \quad a^2 - (n+1)a + a(1-a)\left(\frac{a}{1+a}\right)^n, & \lambda = \mu_i, \end{cases} \\ B_i(n, \text{busy}) &= b \left(\hat{P}_i(\text{busy}) - \hat{P}_i(\text{setup}) \frac{a}{1+a} \right), \end{aligned}$$

where we have used shorthand notation $a = \frac{\lambda}{\delta_i}$ and $b = \frac{\lambda}{\mu_i}$. For any $x = (n, \text{setup})$, where $n \geq 1$,

$$\begin{aligned} H_i(n, \text{setup}) &= \begin{cases} \frac{b}{b(1+a)-a} \left(\frac{b((n+1)-(n+2)b+b^{n+2})}{(b-1)^2} + a \right), & \lambda \neq \mu_i, \\ \frac{1}{2} (n+1)(n+2) + a, & \lambda = \mu_i, \end{cases} \\ B_i(n, \text{setup}) &= b \left(\hat{P}_i(\text{busy}) - \hat{P}_i(\text{setup}) \frac{a}{1+a} \right), \end{aligned}$$

where we have used shorthand notation $a = \frac{\lambda}{\delta_i}$ and $b = \frac{\lambda}{\mu_i}$.

Proof: 1° We first prove that $\pi_i^\emptyset$ is optimal for all

$$\nu \in [0, \nu_i^*(0, \text{off})].$$

Denote $\pi_i^\emptyset$ here briefly by π . By Proposition 1 and the formulas given in Appendix, we get

$$\bar{c}_i^\pi(\nu) = \beta P_i(\text{off}) + \nu.$$

Furthermore, the Howard equations for states (1, setup) and (1, busy) are as follows:

$$\begin{aligned} \bar{c}_i^\pi(\nu) &= h_i + \beta P_i(\text{setup}) + \\ &\quad \delta_i(v_i^\pi(1, \text{busy}; \nu) - v_i^\pi(1, \text{setup}; \nu)) + \nu, \\ \bar{c}_i^\pi(\nu) &= h_i + \beta P_i(\text{busy}) + \\ &\quad \mu_i(v_i^\pi(0, \text{off}; \nu) - v_i^\pi(1, \text{busy}; \nu)) + \nu. \end{aligned}$$

By solving these linear equations, we get

$$\begin{aligned} v_i^\pi(1, \text{setup}; \nu) - v_i^\pi(1, \text{busy}; \nu) &= \frac{1}{\delta_i} (h_i + \beta \hat{P}(\text{setup})), \\ v_i^\pi(1, \text{busy}; \nu) - v_i^\pi(0, \text{off}; \nu) &= \frac{1}{\mu_i} (h_i + \beta \hat{P}(\text{busy})). \end{aligned}$$

Thus,

$$\begin{aligned} \Delta_i^\pi(0, \text{off}; \nu) &= v_i^\pi(1, \text{setup}; \nu) - v_i^\pi(1, \text{busy}; \nu) + \\ &\quad v_i^\pi(1, \text{busy}; \nu) - v_i^\pi(0, \text{off}; \nu) = \\ &\quad h_i \left(\frac{1}{\delta_i} + \frac{1}{\mu_i} \right) + \beta (\hat{P}_i(\text{setup}) \frac{1}{\delta_i} + \hat{P}_i(\text{busy}) \frac{1}{\mu_i}). \end{aligned} \quad (20)$$

Now, the following condition for optimality of π ,

$$\nu < \lambda \Delta_i^\pi(0, \text{off}; \nu)$$

is, by (20), equivalent with

$$\nu < h_i \left(\frac{\lambda}{\delta_i} + \frac{\lambda}{\mu_i} \right) + \beta (\hat{P}_i(\text{setup}) \frac{\lambda}{\delta_i} + \hat{P}_i(\text{busy}) \frac{\lambda}{\mu_i}). \quad (21)$$

Note that the right-hand-side of (21) is positive and equal to $\nu_i^*(0, \text{off})$. Furthermore, by Proposition 2, we deduce that

$$\nu < \lambda \Delta_i^\pi(0, \text{off}; \nu) < \lambda \Delta_i^\pi(x; \nu)$$

for the other states $x \in \{(1, \text{busy}), (1, \text{setup}), (2, \text{busy}), \dots\}$, from which claim 1° follows.

2° Next we prove that $\pi_i^{0, \text{off}}$ is optimal for all

$$\nu \in [\nu_i^*(0, \text{off}), \nu_i^*(1, \text{busy})].$$

Denote $\pi_i^{0, \text{off}}$ here briefly by π . By Proposition 1, we get

$$\bar{c}_i^\pi(\nu) = h_i \frac{D_i^{0, \text{off}}}{G_i^{0, \text{off}}} + \beta \frac{E_i^{0, \text{off}}}{G_i^{0, \text{off}}} + \nu \left(1 - \frac{F_i^{0, \text{off}}}{G_i^{0, \text{off}}} \right).$$

Furthermore, the Howard equation for state (0, off) reads as

$$\bar{c}_i^\pi(\nu) = \beta P_i(\text{off}) + \lambda \Delta_i^\pi(0, \text{off}; \nu),$$

from which we get

$$\begin{aligned} \Delta_i^\pi(0, \text{off}; \nu) &= \frac{1}{\lambda} \left(h_i \frac{D_i^{0, \text{off}}}{G_i^{0, \text{off}}} + \beta \frac{E_i^{0, \text{off}}}{G_i^{0, \text{off}}} + \right. \\ &\quad \left. \nu \left(1 - \frac{F_i^{0, \text{off}}}{G_i^{0, \text{off}}} \right) - \beta P_i(\text{off}) \right). \end{aligned} \quad (22)$$

Now, the following condition for optimality of π ,

$$\nu \geq \lambda \Delta_i^\pi(0, \text{off}; \nu)$$

is, by (22) and the formulas given in Appendix, equivalent with

$$\nu \geq h_i \left(\frac{\lambda}{\delta_i} + \frac{\lambda}{\mu_i} \right) + \beta (\hat{P}_i(\text{setup}) \frac{\lambda}{\delta_i} + \hat{P}_i(\text{busy}) \frac{\lambda}{\mu_i}). \quad (23)$$

Note that the right-hand-side of (23) is equal to $\nu_i^*(0, \text{off})$.

In addition, the Howard equation for state (2, busy) is

$$\bar{c}_i^\pi(\nu) = 2h_i + \beta P_i(\text{busy}) - \mu_i \Delta_i^\pi(1, \text{busy}; \nu) + \nu,$$

from which we get

$$\begin{aligned} \Delta_i^\pi(1, \text{busy}; \nu) &= \frac{1}{\mu_i} \left(2h_i + \beta P_i(\text{busy}) + \nu - h_i \frac{D_i^{0, \text{off}}}{G_i^{0, \text{off}}} \right. \\ &\quad \left. - \beta \frac{E_i^{0, \text{off}}}{G_i^{0, \text{off}}} - \nu \left(1 - \frac{F_i^{0, \text{off}}}{G_i^{0, \text{off}}} \right) \right). \end{aligned} \quad (24)$$

Now, the following condition for optimality of π ,

$$\nu < \lambda \Delta_i^\pi(1, \text{busy}; \nu)$$

is, by (24) and the formulas given in Appendix, equivalent with

$$\begin{aligned} \nu &< h_i \frac{\lambda}{\mu_i} \left(1 + \frac{1 + \frac{\lambda}{\mu_i}}{1 + \frac{\lambda}{\delta_i}} \right) + \\ &\quad \beta \frac{\lambda}{\mu_i} \left(\hat{P}_i(\text{busy}) - \hat{P}_i(\text{setup}) \frac{\frac{\lambda}{\delta_i}}{1 + \frac{\lambda}{\delta_i}} \right). \end{aligned} \quad (25)$$

Note that the right-hand-side of (25) is equal to $\nu_i^*(1, \text{busy})$. Note also that this case is possible only if the right-hand-side of (23) is strictly smaller than right-hand-side of (25). However, it is easy to check that this is guaranteed by our assumption (14). Furthermore, by Proposition 3, we deduce that

$$\nu < \lambda \Delta_i^\pi(1, \text{busy}; \nu) < \lambda \Delta_i^\pi(x; \nu)$$

for the “later” states $x \in \{(1, \text{setup}), (2, \text{busy}), (2, \text{setup}), \dots\}$, from which claim 2° follows.

3° Let then $n \geq 1$. Now we prove that $\pi_i^{n, \text{busy}}$ is optimal for all

$$\nu \in [\nu_i^*(n, \text{busy}), \nu_i^*(n, \text{setup})].$$

Denote $\pi_i^{n, \text{busy}}$ here briefly by π . By Proposition 1, we get

$$\bar{c}_i^\pi(\nu) = h_i \frac{D_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} + \beta \frac{E_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} + \nu \left(1 - \frac{F_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} \right).$$

Furthermore, the Howard equation for state $(n+1, \text{busy})$ is

$$\bar{c}_i^\pi(\nu) = (n+1)h_i + \beta P_i(\text{busy}) - \mu_i \Delta_i^\pi(n, \text{busy}; \nu) + \nu,$$

from which we get

$$\begin{aligned} \Delta_i^\pi(n, \text{busy}; \nu) &= \\ &\frac{1}{\mu_i} \left((n+1)h_i + \beta P_i(\text{busy}) + \nu - h_i \frac{D_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} \right. \\ &\quad \left. - \beta \frac{E_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} - \nu \left(1 - \frac{F_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} \right) \right). \end{aligned} \quad (26)$$

Now, the following condition for optimality of π ,

$$\nu \geq \lambda \Delta_i^\pi(n, \text{busy}; \nu)$$

is, by (26), equivalent with

$$\begin{aligned} \nu &\geq h_i \frac{\lambda}{\mu_i} \frac{(n+1)G_i^{n, \text{busy}} - D_i^{n, \text{busy}}}{G_i^{n, \text{busy}} - \frac{\lambda}{\mu_i} F_i^{n, \text{busy}}} + \\ &\quad \beta \frac{\lambda}{\mu_i} \frac{P_i(\text{busy})G_i^{n, \text{busy}} - E_i^{n, \text{busy}}}{G_i^{n, \text{busy}} - \frac{\lambda}{\mu_i} F_i^{n, \text{busy}}}. \end{aligned} \quad (27)$$

By further applying the formulas given in Appendix, we conclude that the right-hand-side of (27) is equal to $\nu_i^*(n, \text{busy})$.

In addition, the Howard equations for states (n, setup) and $(n+1, \text{setup})$ are as follows:

$$\begin{aligned} \bar{c}_i^\pi(\nu) &= nh_i + \beta P_i(\text{setup}) + \\ &\delta_i(v_i^\pi(n, \text{busy}; \nu) - v_i^\pi(n, \text{setup}; \nu)) + \nu, \\ \bar{c}_i^\pi(\nu) &= (n+1)h_i + \beta P_i(\text{setup}) + \\ &\delta_i(v_i^\pi(n+1, \text{busy}; \nu) - v_i^\pi(n+1, \text{setup}; \nu)) + \nu. \end{aligned}$$

Thus,

$$\begin{aligned} \Delta_i^\pi(n, \text{setup}; \nu) &= \\ &v_i^\pi(n+1, \text{setup}; \nu) - v_i^\pi(n+1, \text{busy}; \nu) + \\ &\Delta_i^\pi(n, \text{busy}; \nu) + v_i^\pi(n, \text{busy}; \nu) - v_i^\pi(n, \text{setup}; \nu) = \\ &\Delta_i^\pi(n, \text{busy}; \nu) + \frac{h_i}{\delta_i}, \end{aligned}$$

which gives

$$\begin{aligned} \Delta_i^\pi(n, \text{setup}; \nu) &= \\ &\frac{1}{\mu_i} \left((n+1 + \frac{\mu_i}{\delta_i}) h_i + \beta P_i(\text{busy}) + \nu - h_i \frac{D_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} \right. \\ &\quad \left. - \beta \frac{E_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} - \nu \left(1 - \frac{F_i^{n, \text{busy}}}{G_i^{n, \text{busy}}} \right) \right). \end{aligned} \quad (28)$$

Now, the following condition for optimality of π ,

$$\nu < \lambda \Delta_i^\pi(n, \text{setup}; \nu)$$

is, by (28), equivalent with

$$\begin{aligned} \nu < h_i \frac{\lambda}{\mu_i} \frac{(n+1+\frac{\mu_i}{\delta_i})G_i^{n,\text{busy}} - D_i^{n,\text{busy}}}{G_i^{n,\text{busy}} - \frac{\lambda}{\mu_i} F_i^{n,\text{busy}}} + \\ \beta \frac{\lambda}{\mu_i} \frac{G_i^{n,\text{busy}} P_i(\text{busy}) - E_i^{n,\text{busy}}}{G_i^{n,\text{busy}} - F_i^{n,\text{busy}} \frac{\lambda}{\mu_i}}. \end{aligned} \quad (29)$$

By further applying the formulas given in Appendix, we conclude that the right-hand-side of (29) is equal to $\nu_i^*(n, \text{setup})$. Note that this case is possible only if the right-hand-side of (27) is strictly smaller than right-hand-side of (29), which, however, is clear from these formulas, since $n+1 < n+1+\frac{\mu_i}{\delta_i}$.

Furthermore, by Proposition 4, we deduce that

$$\nu < \lambda \Delta_i^\pi(n, \text{setup}; \nu) < \lambda \Delta_i^\pi(x; \nu)$$

for the “later” states $x \in \{(n+1, \text{busy}), (n+1, \text{setup}), (n+2, \text{busy}), \dots\}$. Moreover, it is possible to show that

$$\nu \geq \lambda \Delta_i^\pi(n, \text{busy}; \nu) \geq \lambda \Delta_i^\pi(x; \nu)$$

for the “prior” states $x \in \{(0, \text{off}), (1, \text{busy}), \dots, (n-1, \text{setup})\}$, from which claim 3° follows.

4° Let again $n \geq 1$. We finally prove that $\pi_i^{n,\text{setup}}$ is optimal for all

$$\nu \in [\nu_i^*(n, \text{setup}), \nu_i^*(n+1, \text{busy})].$$

Denote $\pi_i^{n,\text{setup}}$ here briefly by π . By Proposition 1, we get

$$\bar{c}_i^\pi(\nu) = h_i \frac{D_i^{n,\text{setup}}}{G_i^{n,\text{setup}}} + \beta \frac{E_i^{n,\text{setup}}}{G_i^{n,\text{setup}}} + \nu \left(1 - \frac{F_i^{n,\text{setup}}}{G_i^{n,\text{setup}}}\right).$$

Furthermore, the Howard equations for states (n, setup) , $(n+1, \text{busy})$, and $(n+1, \text{setup})$ are given by

$$\begin{aligned} \bar{c}_i^\pi(\nu) &= nh_i + \beta P_i(\text{setup}) + \\ &\delta_i(v_i^\pi(n, \text{busy}; \nu) - v_i^\pi(n, \text{setup}; \nu)) + \lambda \Delta_i^\pi(n, \text{setup}; \nu), \\ \bar{c}_i^\pi(\nu) &= (n+1)h_i + \beta P_i(\text{busy}) - \mu_i \Delta_i^\pi(n, \text{busy}; \nu) + \nu, \\ \bar{c}_i^\pi(\nu) &= (n+1)h_i + \beta P_i(\text{setup}) + \\ &\delta_i(v_i^\pi(n+1, \text{busy}; \nu) - v_i^\pi(n+1, \text{setup}; \nu)) + \nu, \end{aligned}$$

from which we get

$$\begin{aligned} v_i^\pi(n, \text{busy}; \nu) - v_i^\pi(n, \text{setup}; \nu) &= \\ &-\frac{1}{\delta_i}(nh_i + \beta \hat{P}(\text{setup}) + \lambda \Delta_i^\pi(n, \text{setup}; \nu) - \bar{c}_i^\pi(\nu)), \\ \Delta_i^\pi(n, \text{busy}; \nu) &= \\ &\frac{1}{\mu_i}((n+1)h_i + \beta \hat{P}(\text{busy}) + \nu - \bar{c}_i^\pi(\nu)), \\ v_i^\pi(n+1, \text{setup}; \nu) - v_i^\pi(n+1, \text{busy}; \nu) &= \\ &\frac{1}{\delta_i}((n+1)h_i + \beta \hat{P}(\text{setup}) + \nu - \bar{c}_i^\pi(\nu)). \end{aligned}$$

Thus,

$$\begin{aligned} \Delta_i^\pi(n, \text{setup}; \nu) &= \\ v_i^\pi(n+1, \text{setup}; \nu) - v_i^\pi(n+1, \text{busy}; \nu) + \\ \Delta_i^\pi(n, \text{busy}; \nu) + v_i^\pi(n, \text{busy}; \nu) - v_i^\pi(n, \text{setup}; \nu) &= \\ \Delta_i^\pi(n, \text{busy}; \nu) + \frac{1}{\delta_i}(h_i + \nu - \lambda \Delta_i^\pi(n, \text{setup}; \nu)), \end{aligned}$$

which gives

$$\begin{aligned} \Delta_i^\pi(n, \text{setup}; \nu) &= \\ \frac{1}{1+\frac{\lambda}{\delta_i}} \left[\frac{1}{\mu_i} \left((n+1)h_i + \beta \hat{P}(\text{busy}) + \nu - h_i \frac{D_i^{n,\text{setup}}}{G_i^{n,\text{setup}}} \right. \right. \\ &\quad \left. \left. - \beta \frac{E_i^{n,\text{setup}}}{G_i^{n,\text{setup}}} - \nu \left(1 - \frac{F_i^{n,\text{setup}}}{G_i^{n,\text{setup}}}\right) \right) + \frac{1}{\delta_i}(h_i + \nu) \right]. \end{aligned} \quad (30)$$

Now, the following condition for optimality of π ,

$$\nu \geq \lambda \Delta_i^\pi(n, \text{setup}; \nu)$$

is, by (30), equivalent with

$$\begin{aligned} \nu \geq h_i \frac{\lambda}{\mu_i} \frac{(n+1+\frac{\mu_i}{\delta_i})G_i^{n,\text{setup}} - D_i^{n,\text{setup}}}{G_i^{n,\text{setup}} - \frac{\lambda}{\mu_i} F_i^{n,\text{setup}}} + \\ \beta \frac{\lambda}{\mu_i} \frac{P_i(\text{busy})G_i^{n,\text{setup}} - E_i^{n,\text{setup}}}{G_i^{n,\text{setup}} - \frac{\lambda}{\mu_i} F_i^{n,\text{setup}}}. \end{aligned} \quad (31)$$

By further applying the formulas given in Appendix, we conclude that the right-hand-side of (31) is equal to $\nu_i^*(n, \text{setup})$.

In addition, the Howard equation for state $(n+2, \text{busy})$ is

$$\bar{c}_i^\pi(\nu) = (n+2)h_i + \beta P_i(\text{busy}) - \mu_i \Delta_i^\pi(n+1, \text{busy}; \nu) + \nu,$$

from which we get

$$\begin{aligned} \Delta_i^\pi(n+1, \text{busy}; \nu) &= \\ \frac{1}{\mu_i} \left((n+2)h_i + \beta P_i(\text{busy}) + \nu - h_i \frac{D_i^{n,\text{setup}}}{G_i^{n,\text{setup}}} \right. \\ &\quad \left. - \beta \frac{E_i^{n,\text{setup}}}{G_i^{n,\text{setup}}} - \nu \left(1 - \frac{F_i^{n,\text{setup}}}{G_i^{n,\text{setup}}}\right) \right). \end{aligned} \quad (32)$$

Now, the following condition for optimality of π ,

$$\nu < \lambda \Delta_i^\pi(n+1, \text{busy}; \nu)$$

is, by (32), equivalent with for this case:

$$\begin{aligned} \nu < h_i \frac{\lambda}{\mu_i} \frac{(n+2)G_i^{n,\text{setup}} - D_i^{n,\text{setup}}}{G_i^{n,\text{setup}} - \frac{\lambda}{\mu_i} F_i^{n,\text{setup}}} + \\ \beta \frac{\lambda}{\mu_i} \frac{P_i(\text{busy})G_i^{n,\text{setup}} - E_i^{n,\text{setup}}}{G_i^{n,\text{setup}} - \frac{\lambda}{\mu_i} F_i^{n,\text{setup}}}. \end{aligned} \quad (33)$$

By further applying the formulas given in Appendix, we conclude that the right-hand-side of (33) is equal to $\nu_i^*(n+1, \text{busy})$.

Note that this case is possible only if the right-hand-side of (31) is strictly smaller than right-hand-side of (33), which, however, is easily verified from these formulas, since $n+1+\frac{\mu_i}{\delta_i} < n+2$ by our assumption (14).

Furthermore, by Proposition 5, we deduce that

$$\nu < \lambda \Delta_i^\pi(n+1, \text{busy}; \nu) < \lambda \Delta_i^\pi(x; \nu)$$

for the “later” states $x \in \{(n+1, \text{setup}), (n+2, \text{busy}), (n+2, \text{setup}), \dots\}$. Moreover, it is possible to show that

$$\nu \geq \lambda \Delta_i^\pi(n, \text{setup}; \nu) \geq \lambda \Delta_i^\pi(x; \nu)$$

for the “prior” states $x \in \{(0, \text{off}), (1, \text{busy}), \dots, (n, \text{busy})\}$, from which claim 4° follows.

By combining 1°–4°, we conclude that the claim is true. ■

From Proposition 2, we see that our condition (14) comes from the requirement that the value function differences $\Delta_i^\pi(1, \text{busy}; \nu)$ and $\Delta_i^\pi(0, \text{off}; \nu)$ for the rudimentary threshold policy $\pi = \pi_i^\emptyset$, which rejects all the incoming jobs, satisfy

$$\Delta_i^\pi(1, \text{busy}; \nu) > \Delta_i^\pi(0, \text{off}; \nu).$$

This implies that when ν is continuously increased starting from zero, the new jobs are first dispatched to server i in state $(0, \text{off})$, laying the ground for the order of states given in (15).

Note that (14) is a sufficient condition for indexability but not necessary. However, it seems that, for some other condition, a different ordering of states is implied.

IV. ENERGY-AWARE WHITTLE INDEX POLICY

Now we return to the original dispatching problem described in Section II, where we have K servers and a strict dispatching condition (3). Based on the results given in Section III, we introduce the following energy-aware index policy.

Definition 1: For any InstantOff server i with state $x_i = (n_i, z_i)$, we define index

$$\nu_i^{\text{EW}}(x_i) = \nu_i^*(x_i),$$

where $\nu_i^*(x)$ is defined by (19) in Theorem 1. The dispatching rule that at every time t chooses the server with the lowest index $\nu_i^{\text{EW}}(x_i)$ is called the *Energy-aware Whittle index policy* (EW) for the original dispatching problem.

V. NUMERICAL RESULTS

Here we illustrate the performance of the Whittle-index based EW policy and compare it against several reference policies to gain insight. The first reference policy is the static LB (load balancing) policy, which is a probabilistic policy where the probability of dispatching the arrival to the i^{th} queue is proportional to μ_i . Thus, the LB policy equalizes the load in each queue, and it is independent of the dynamic state of the system being still maximally stable. The JSQ (Join-the-Shortest-Queue) policy is a dynamic policy that dispatches the arrival to the queue with the smallest number of jobs, and breaks ties randomly. The FPI (First Policy Iteration) policy is also a dynamic policy, derived for our dispatching problem in [19], and it is based on the well-known theory of MDPs and the policy iteration algorithm [27]. It uses the LB policy as the initial policy and performs the optimal dispatching action for an arrival assuming that the initial policy is used after that. Similar to our proposed EW policy, the FPI policy is near-optimal and we are interested in seeing which one is indeed closer to the optimal. Finally, we apply the policy iteration algorithm to numerically solve the optimal policy, whenever possible.

In our examples, the results for the policies JSQ, FPI and EW have been produced by using discrete-event simulations. For each combination of the parameters, the results are based on simulation runs with $2 \cdot 10^6$ arrivals. The LB policy can be easily evaluated numerically, since under that policy each queue behaves independently from each other as an M/G/1 queue with setup delays, see [19]. Finally, the results are

shown as a function of the total load of the system, given by $\rho = \lambda / \sum_i \mu_i$. To vary the load ρ , we fix the μ_i and vary λ . Also, in our examples for the power settings, we assume $P_i(\text{busy}) = P_i(\text{setup})$ and also $P_i(\text{off}) = 0$, for all i .

In our first example, we consider a small system with 2 queues. The parameters are the following: $\{\mu_1, \mu_2\} = \{1, 4\}$, $\{\delta_1, \delta_2\} = \{4, 17\}$, $\{P_1(\text{busy}), P_2(\text{busy})\} = \{200, 300\}\text{W}$ and $\beta = 0.001$. The parameters satisfy our assumption (14). For this small system, the state space of the 4 dimensional process remains moderate and we are able to apply the policy iteration algorithm to numerically solve the optimal policy minimizing the mean total costs (5) in a truncated state space. Here the truncation has been done at 30 jobs in each queue, which is sufficiently high to allow a reasonably accurate estimation of the optimal policy even at relatively high values of the load ρ . For each load, the policy iteration has been performed for 10 iterations, which yielded the optimal costs with at least 5 digit accuracy. The results are shown in Figure 1, which depicts the ratio of the mean number of jobs (left panel), the mean power (middle panel) and the mean total cost (right panel) to the optimal policy as a function of the load ρ for different policies (LB, JSQ, FPI and EW).

Consider first the performance ratio (left panel). The static LB policy clearly does not perform very well, and gets worse as load increases relative to the optimal policy. On the other hand the JSQ policy gets better and better the higher the load, which can be explained by the fact that at very high load all queues are active all the time and the energy-aware features are not affecting the behavior that much. The servers are heterogeneous, but JSQ also indirectly takes care of this as the queue in the faster server is typically shorter than in the slower one. The FPI and EW policies are both clearly near-optimal, i.e., the ratio is close to 1. However, the FPI policy is less optimal than the EW policy. Then considering the results for the power ratio (middle panel), it can be observed that the non-energy-aware JSQ policy is the worst, being even worse than the static LB policy. Both LB and JSQ are performing poorly at low loads but become better as load increases, since at higher loads both servers are on all the time anyway. However, the FPI policy is here even better with respect power consumption than the optimal policy, and it is also better than EW, which remains very close to optimal until $\rho = 0.6$ but then becomes marginally better than the optimal policy. Finally, by looking at the total costs (right panel) we see that they are close to the ones for the performance part as the weight $\beta = 0.001$ is quite small. In summary, our proposed EW policy is performing systematically better than FPI and it is overall very close to the optimal policy; in fact it is indistinguishable from the optimal until load $\rho = 0.6$, while the performance of the FPI policy starts deviating from the optimal already at low load, reaching a deviation of approximately 10% at high load.

Next we consider a larger 4-server system with the following parameters satisfying (14): $\{\mu_1, \mu_2, \mu_3, \mu_4\} = \{1, 2, 3, 4\}$, $\{\delta_1, \delta_2, \delta_3, \delta_4\} = \{4, 9, 16, 25\}$, $P_i(\text{busy}) = (100 + i \cdot 100)\text{W}$, $i = 1, \dots, 4$, and $\beta = 0.001$ or $\beta = 0.01$. Now the state space of the process is too large to allow for solving numerically

Fig. 1. Ratio of mean number of jobs (left panel), mean power (middle panel) and mean total cost (right panel) to the corresponding quantities of the optimal policy as a function of the load ρ for different policies (LB, JSQ, FPI and EW) with 2 servers.

the optimal policy. The results are shown in Figure 2, which shows the mean total cost as a function of the load for different policies (LB, JSQ, FPI, EQ) for $\beta = 0.001$ (upper panel) and $\beta = 0.01$ (lower panel). We consider first the case with $\beta = 0.001$ (upper panel). Overall, the mean cost rate starts increasing quickly for all policies as load increases towards 1, representing the stability limit. Again LB performs worst and all dynamic policies are significantly better. At low load, JSQ is worse than FPI and EW, but as load increases JSQ becomes better. The EW policy remains throughout better than FPI, similarly as in Figure 1. By looking at the performance and the power separately (not shown here due to lack of space), the gain for the EW policy over the FPI policy could be seen to follow from the performance part and it dominates as $\beta = 0.001$ is small. By increasing $\beta = 0.01$, see Figure 2 (lower panel), we observe that the gap between FPI and EW diminishes, because the lower power consumption of the FPI policy reduces the performance benefits of the EW policy.

Fig. 2. Mean total cost for different policies (LB, JSQ, FPI and EW) as a function of the load ρ for 4 servers.

VI. CONCLUSIONS

We have considered the energy-aware dispatching problem in a system consisting of parallel M/M/1 queues with InstantOff servers. Such servers go to sleep after becoming empty to save energy, but activation of the server after sleep incurs an additional setup delay penalty. The costs in our system model consist of linear holding costs and power consumption costs. The performance-energy trade-off is characterized as a weighted sum of these.

To optimize the trade-off we have applied the Whittle index approach, which is based on a certain relaxation of the original intractable dispatching problem, and results in a separable problem, where each queue is considered in isolation. As our main result, we have proved, under a certain condition, the indexability property, and we have also derived the explicit form of the Whittle index. The proof is technically challenging, as each queue is described by a two-dimensional Markov process, representing the M/M/1 queue with setup, for which indexability results are not available in the literature. As demonstrated by our numerical results, the energy-aware Whittle index policy is able to perform very close to the numerically solved optimal policy and outperforms all considered reference policies.

In this paper, indexability is proved and the corresponding Whittle index is derived under a certain condition related to the system parameters. Future research includes relaxing this condition, although it may be challenging.

REFERENCES

- [1] M. Harchol-Balter, *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*. Cambridge University Press, 2013.
- [2] W. Winston, "Optimality of the shortest line discipline," *Journal of Applied Probability*, vol. 14, pp. 181–189, 1977.
- [3] R. Weber, "On the optimal assignment of customers to parallel servers," *Journal of Applied Probability*, vol. 15, pp. 406–413, 1978.
- [4] P. Whittle, *Optimal Control: Basics and Beyond*. Wiley, 1996.
- [5] P. Ansell, K. Glazebrook, and C. Kirkbride, "Generalised join the shortest queue policies for the dynamic routing of jobs to multiclass queues," *Journal of the Operational Research Society*, vol. 54, pp. 379–389, 2003.
- [6] N. Argon, L. Ding, K. Glazebrook, and S. Ziya, "Dynamic routing of customers with general delay costs in a multiserver queueing system," *Probability in the Engineering and Informational Sciences*, vol. 23, pp. 175–203, 2009.

- [7] E. Hytiä, A. Penttinen, and S. Aalto, "Size- and state-aware dispatching problem with queue-specific job sizes," *European Journal of Operational Research*, vol. 217, pp. 357–370, 2012.
- [8] E. Hytiä, S. Aalto, A. Penttinen, and J. Virtamo, "On the value function of the M/G/1 FCFS and LCFS queues," *Journal of Applied Probability*, vol. 49, pp. 1052–1071, 2012.
- [9] E. Hytiä, R. Richter, and J. Virtamo, "Meeting soft deadlines in single- and multi-server systems," in *Proc. of ITC 28*, September 2016, pp. 166–174.
- [10] A. Gandhi, V. Gupta, M. Harchol-Balter, and M. Kozuch, "Optimality analysis of energy-performance trade-off for server farm management," *Performance Evaluation*, vol. 67, pp. 1155–1171, 2010.
- [11] V. Maccio and D. Down, "On optimal policies for energy-aware servers," *Performance Evaluation*, vol. 90, pp. 36–52, 2015.
- [12] M. Gebrehiwot, S. Aalto, and P. Lassila, "Optimal energy-aware control policies for FIFO servers," *Performance Evaluation*, vol. 103, pp. 41–59, 2016.
- [13] —, "Energy-performance trade-off for processor sharing queues with setup delay," *Operations Research Letters*, vol. 44, pp. 101–106, 2016.
- [14] —, "Energy-aware SRPT server with batch arrivals: Analysis and optimization," *Performance Evaluation*, vol. 115, pp. 92–107, 2017.
- [15] E. Hytiä, R. Richter, and S. Aalto, "Task assignment in a heterogeneous server farm with switching delays and general energy-aware cost structure," *Performance Evaluation*, vol. 75–76, pp. 17–35, 2014.
- [16] —, "Energy-aware job assignment in server farms with setup delays under LCFS and PS," in *Proc. of ITC 26*, September 2014.
- [17] E. Hytiä and R. Richter, "Fairness through linearly increasing holding costs in systems of parallel servers with setup delays," in *Proc. of ITC 27*, September 2015, pp. 143–151.
- [18] E. Hytiä, R. Richter, J. Virtamo, and L. Viitasari, "Value (generating) functions for the M^X/G/1 queue," in *Proc. of ITC 29*, September 2017, pp. 232–240.
- [19] M. Gebrehiwot, S. Aalto, and P. Lassila, "Near-optimal policies for energy-aware task assignment in server farms," in *Proc. of CCGrid*, May 2017, pp. 1017–1026.
- [20] P. Whittle, "Restless bandits: activity allocation in a changing world," *Journal of Applied Probability*, vol. 25A, pp. 287–298, 1988.
- [21] U. Ayesta, M. Erasquin, and P. Jacko, "A modeling framework for optimizing the flow-level scheduling with time-varying channels," *Performance Evaluation*, vol. 67, pp. 1014–1029, 2010.
- [22] P. Jacko, "Value of information in optimal flow-level scheduling of users with Markovian time-varying channels," *Performance Evaluation*, vol. 68, pp. 1022–1036, 2011.
- [23] F. Cecchi and P. Jacko, "Nearly-optimal scheduling of users with Markovian time-varying transmission rates," *Performance Evaluation*, vol. 99–100, pp. 16–36, 2016.
- [24] S. Aalto, P. Lassila, and P. Osti, "Whittle index approach to size-aware scheduling for time-varying channels with multiple states," *Queueing Systems*, vol. 83, pp. 195–225, 2016.
- [25] J. Niño-Mora, "Dynamic allocation indices for restless projects and queueing admission control: A polyhedral approach," *Mathematical Programming*, vol. 93, pp. 361–413, 2002.
- [26] M. Larrañaga, U. Ayesta, and I. Verloop, "Dynamic control of birth-and-death restless bandits: Application to resource-allocation problems," *IEEE/ACM Transactions on Networking*, vol. 24, pp. 3812–3825, 2016.
- [27] M. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 2005.

APPENDIX

Below we give the definitions for the constants D_i^π , E_i^π , F_i^π , and G_i^π introduced in Proposition 1.

For $\pi = \pi_i^\emptyset$, we define

$$D_i^\emptyset = 0, \quad E_i^\emptyset = P_i(\text{off}), \quad F_i^\emptyset = 0, \quad G_i^\emptyset = 1. \quad (34)$$

For $\pi = \pi_i^{0,\text{off}}$, we define

$$\begin{aligned} D_i^{0,\text{off}} &= \frac{\lambda}{\delta_i} + \frac{\lambda}{\mu_i}, \\ E_i^{0,\text{off}} &= P_i(\text{off}) + P_i(\text{setup}) \frac{\lambda}{\delta_i} + P_i(\text{busy}) \frac{\lambda}{\mu_i}, \\ F_i^{0,\text{off}} &= 1, \\ G_i^{0,\text{off}} &= 1 + \frac{\lambda}{\delta_i} + \frac{\lambda}{\mu_i}. \end{aligned} \quad (35)$$

Next, for any $\pi = \pi_i^{n,\text{busy}}$, where $n \geq 1$, we define (assuming first that $\lambda \neq \mu_i$)

$$\begin{aligned} D_i^{n,\text{busy}} &= \frac{((n+1)b+a)(a-b)a^n}{b(1+a)-a} + \frac{(b^{n+3}((n+1)b-(n+2))+b^2-(b-1)^2a^2)(1+a)^n}{(b-1)^2(b(1+a)-a)}, \\ E_i^{n,\text{busy}} &= P_i(\text{off})(1+a)^{n-1} + P_i(\text{setup})a(1+a)^{n-1} + P_i(\text{busy}) \left(\frac{b(a-b)a^n}{b(1+a)-a} + \frac{b(b^{n+2}-b-(b-1)a)(1+a)^n}{(b-1)(b(1+a)-a)} \right), \\ F_i^{n,\text{busy}} &= \frac{(a-b)a^n}{b(1+a)-a} + \frac{(b^{n+2}-b-(b-1)a)(1+a)^n}{(b-1)(b(1+a)-a)}, \\ G_i^{n,\text{busy}} &= \frac{b(a-b)a^n}{b(1+a)-a} + \frac{(b^{n+3}-b-(b-1)a)(1+a)^n}{(b-1)(b(1+a)-a)}, \end{aligned} \quad (36)$$

where we have used shorthand notation $a = \frac{\lambda}{\delta_i}$ and $b = \frac{\lambda}{\mu_i}$. However, if $\lambda = \mu_i$, we define

$$\begin{aligned} D_i^{n,\text{busy}} &= (n+1+a)(a-1)a^n + \left(\frac{1}{2}(n+1)(n+2)-a^2 \right)(1+a)^n, \\ E_i^{n,\text{busy}} &= P_i(\text{off})(1+a)^{n-1} + P_i(\text{setup})a(1+a)^{n-1} + P_i(\text{busy})((a-1)a^n + (n+1-a)(1+a)^n), \\ F_i^{n,\text{busy}} &= (a-1)a^n + (n+1-a)(1+a)^n, \\ G_i^{n,\text{busy}} &= (a-1)a^n + (n+2-a)(1+a)^n, \end{aligned} \quad (37)$$

where we have used shorthand notation $a = \frac{\lambda}{\delta_i}$.

Finally, for any $\pi = \pi_i^{n,\text{setup}}$, where $n \geq 1$, we define (assuming first that $\lambda \neq \mu_i$)

$$\begin{aligned} D_i^{n,\text{setup}} &= \frac{((n+1)b+a)a^{n+2}}{b(1+a)-a} + \frac{(b^{n+3}((n+1)b-(n+2))+b^2-(b-1)^2a^2)(1+a)^{n+1}}{(b-1)^2(b(1+a)-a)}, \\ E_i^{n,\text{setup}} &= P_i(\text{off})(1+a)^n + P_i(\text{setup})a(1+a)^n + P_i(\text{busy}) \left(\frac{ba^{n+2}}{b(1+a)-a} + \frac{b(b^{n+2}-b-(b-1)a)(1+a)^{n+1}}{(b-1)(b(1+a)-a)} \right), \\ F_i^{n,\text{setup}} &= \frac{a^{n+2}}{b(1+a)-a} + \frac{(b^{n+2}-b-(b-1)a)(1+a)^{n+1}}{(b-1)(b(1+a)-a)}, \\ G_i^{n,\text{setup}} &= \frac{ba^{n+2}}{b(1+a)-a} + \frac{(b^{n+3}-b-(b-1)a)(1+a)^{n+1}}{(b-1)(b(1+a)-a)}, \end{aligned} \quad (38)$$

where we have used shorthand notation $a = \frac{\lambda}{\delta_i}$ and $b = \frac{\lambda}{\mu_i}$. However, if $\lambda = \mu_i$, we define

$$\begin{aligned} D_i^{n,\text{setup}} &= (n+1+a)a^{n+2} + \left(\frac{1}{2}(n+1)(n+2)-a^2 \right)(1+a)^{n+1}, \\ E_i^{n,\text{setup}} &= P_i(\text{off})(1+a)^n + P_i(\text{setup})a(1+a)^n + P_i(\text{busy})(a^{n+2} + (n+1-a)(1+a)^{n+1}), \\ F_i^{n,\text{setup}} &= a^{n+2} + (n+1-a)(1+a)^{n+1}, \\ G_i^{n,\text{setup}} &= a^{n+2} + (n+2-a)(1+a)^{n+1}, \end{aligned} \quad (39)$$

where we have used shorthand notation $a = \frac{\lambda}{\delta_i}$.