

Dynamic Bandwidth Allocation for Low Power Devices With Random Connectivity

Navid Ehsan and Mingyan Liu

Abstract—In this paper we consider the bandwidth allocation problem where multiple low power wireless devices share a common time-slotted channel for transmitting to a single server. Due to energy constraints, these devices alternate between common active and inactive periods, the former typically much smaller than the latter. At the beginning of each active period the server decides which user(s) can access the common channel. This decision is based on the knowledge of the current backlog and connectivity of each queue. In each time slot an active user may or may not be connected to the server. If a user is connected to the server, it can transmit with a certain success probability. Arrivals are arbitrary and there is a cost for holding a packet in the queue. Different queues have different packet holding costs leading to differentiated services. We consider the problem of minimizing the total discounted cost over a finite or infinite horizon and provide sufficient conditions under which a greedy policy is optimal. We consider two connectivity models: (1) there is no information about connectivity statistics, and (2) connectivity probability is independent from one time slot to the other (memoryless channel). We show that in each of these cases it is optimal to serve the user with the highest one step reward (smallest one step cost) if this gain is sufficiently larger than that from serving the other users. The sufficient condition is shown to be asymptotically tight in special cases. We then use numerical examples to study the performance of the greedy policy as a function of the duty cycle and the length of the active period. This helps us to better understand and model the tradeoff between increasing the lifetime and decreasing the packet delay in such systems.

Index Terms—Resource allocation, low-power devices, stochastic systems, optimal control.

I. INTRODUCTION

In this paper we consider the problem of allocating bandwidth to multiple users that share a common channel for transmitting to a single server. Users are low power wireless devices, e.g., wireless sensors. In order to conserve energy, they are heavily duty cycled, i.e., they alternate between on/active and off/inactive periods (with the latter typically much larger than the former). All transmissions and receptions occur during the active period and the radio transceivers are turned off during the inactive period. Users' active/inactive schedules are synchronized, in that they wake up and utilize the channel during a common active period and sleep during a common inactive period. An example of such system is the IEEE 802.15.4 (also known as the ZigBee standard for indoor wireless communications) specification, under which devices can be configured to be off for up to

4 minutes at a time, while active for a small fraction of a minute.

In this study we assume that the users share the channel during the common active period via a dynamic TDMA schedule. Specifically, we assume that each active period consists of one or more transmission slots which are allocated to the users by the server. We consider a two state channel model, where a user can transmit a packet with a certain success probability in a slot if it has a good channel. If it has a bad channel then the success probability is zero (i.e., it cannot transmit). At the beginning of each active period, the server is assumed to know the current backlog in the user queues and the channel state of each user. With this information the server allocates the slots and the allocation is announced prior to the transmissions from the users.

In order to measure the performance of the system we consider a cost for keeping packets in the system. This cost is charged for each time slot a packet remains in the queue. Each user may have a different holding cost, which allows us to study differentiated services. The goal is to minimize the total discounted holding cost over a finite or an infinite horizon.

We define a greedy policy that minimizes only the one step cost. This policy is not always optimal when minimizing the cost over a long period of time. On the other hand this policy has a very simple form and only requires the current backlog and connectivity information of each queue to make the allocation decision. Due to its simplicity it is desirable to see under what conditions this policy is optimal and more generally how it performs when these conditions are not satisfied. Furthermore, we are also interested in understanding the relationship between its performance and parameters like the duty cycle and the frequency at which the system duty cycles.

In this paper we first derive the sufficient conditions under which this greedy policy is optimal, and show that the sufficient condition is tight in certain cases. Via numerical examples we study its characteristics under scenarios that do not necessarily satisfy those sufficient conditions. We also show its performance when the system operates under different duty cycles.

Similar resource allocation problems have been extensively studied in the literature. Here we review the ones most relevant to our study in this paper. In the case of linear holding cost and constant success probabilities, the optimality of the greedy policy (also known as the $c\mu$ rule) has been shown in many cases using different arguments. [1] used a dynamic programming argument to show the

N. Ehsan and M. Liu are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor nehsan@umich.edu, mingyan@eecs.umich.edu

optimality of the $c\mu$ rule when $N = 2$ for infinite horizon. [2] used an interchange argument to further show its optimality for $N \geq 2$ in both finite and infinite horizons and for arbitrary arrival processes. Later [3] showed that the region corresponding to the admissible policies is a polymatroid. Using this argument and the results from [4] they proved the optimality of the $c\mu$ rule (also see [5], [6]). In all these scenarios the channel state (connectivity) is fixed and does not change over time.

[7], [8], [9] considered the server allocation problem to multiple queues with varying connectivity but of the same service class. Each of these studies determined policies that maximize throughput over an infinite horizon. In particular, [7] derived the sufficient condition for stability and showed that the Longest Connected Queue (LCQ) policy stabilizes the system if system can be stabilized. [10] further considered a similar problem but with differentiated service classes where different queues have different holding costs. We use some of the ideas in this paper when deriving sufficient conditions for the optimality of the greedy policy. [11], [12] studied the stability of power allocation policies. [13] studied the problem of optimal routing and server allocation for two queues and proved that the optimal policy is of the threshold type, under linear cost functions and uncontrolled arrivals.

Also relevant to the optimal bandwidth allocation problem is the *restless bandit problem*. This class of problems were studied in numerous studies, see for example [14], [15], [16], [17]. An optimal solution for the general restless bandit problem is not known. The problem considered in this paper can be viewed as a special case of the restless bandit problem.

The rest of the paper is organized as follows. In the next section we explain the system model and state the optimization problem. We also define the greedy policy that has a very simple form and present it in terms of an index policy. In section III we provide sufficient conditions for the greedy policy to be optimal. In section IV we study the performance of the greedy policy under various system parameters and show the tradeoff between increasing the lifetime and decreasing the queue cost. Section V concludes the paper.

II. PROBLEM FORMULATION

A. System Description

In this section we describe the model abstraction we adopted for this study. This model is primarily derived from the IEEE 802.15.4 specification as mentioned earlier.

Consider N users sharing a common channel to send packets to a single server. Time is slotted and indexed by $t = 0, 1, 2, \dots$. Users alternate between on/active and off/inactive periods, and the *duty cycle* is defined as the fraction of time a user is on/active. An active period is M slots in length ($M \geq 1$), and an inactive period is $M \cdot (L - 1)$ slots in length ($L \geq 1$), resulting in a duty cycle of $1/L$. For simplicity, L is assumed to be an integer. In words, during

each cycle/period, the system is active for M slots (e.g., from $t = 0$ to $t = M - 1$ for the first cycle) and then goes to sleep for $(L - 1) \cdot M$ slots. Then it becomes active again at the beginning of the LM -th slot and the same process repeats.

The synchronization of users to ensure they adhere to the same active schedule is typically maintained via a beacon sent by the server right at the end of each cycle. That is, a user will wake up right before an active period, wait for the beacon, and resynchronize (e.g., to adjust clock drift, etc.). Since the beacon is typically very short compared to the slot or the cycle, we will ignore its duration. We assume that at the beginning of each active period, the server has the backlog and connectivity information (defined below) of each user in the system. This is a simplification. In reality, this information may be communicated via certain designated mini-slots at the beginning of each active period (between the beacon and the first slot). The allocation decision is then made by the server and announced in a second beacon. For simplicity, the duration of these mini-slots are not considered in our formulation, although this does not affect the applicability of our analysis or results. The above system is shown in Figure 1.

During each active time slot a user may or may not be connected to the server. If the user is not connected to the server, then it cannot transmit. Let $q_{i,t}$ denote the connectivity of user i at time t . If user i is connected to the server at time t then $q_{i,t} = 1$, otherwise it is zero. If user i is connected to the server and the slot is allocated to it, then it transmits a packet successfully with probability p_i . These success probabilities are assumed known by the server. At the beginning of an active period, say slot t , the server observes the queue-size of all the queues, denoted by \mathbf{b}_t , and their connectivity \mathbf{q}_t . The server uses this information to allocate the M slots among users and announces the allocation decision. This is achieved via the mini-slots and the second beacon as described above. The users subsequently follow the allocation decision and transmit in their designated slots (if they are assigned any). The whole system then goes to sleep for $(L - 1)M$ time slots. At time $t = LM$ the cycle repeats (ignoring the time spent by beacons and mini-slots).

We assume that a packet in queue i incurs a holding cost of c_i for every slot it remains in the queue. The cost is collected at the end of each time slot. For instance the cost of time slot t (time interval $[t, t + 1]$) of queue i is equal to $c_i b_{i,[t+1]^-}$ where $b_{i,[t+1]^-}$ is the backlog of queue i right before time $t + 1$. The objective is to find an allocation policy π that minimizes the following cost function:

$$\begin{aligned} J_T^\pi &= E^\pi[C|\mathcal{F}_0], \\ C &= \sum_{t=0}^{T-1} \beta^t \sum_{i=1}^N c_i b_{i,[t+1]^-}, \end{aligned} \quad (1)$$

where \mathcal{F}_0 summarizes all the information available at time $t = 0$, and $\beta < 1$ is the discount factor.




Fig. 1. The low duty cycle dynamics

Consider a greedy policy π^* defined as follows for a single slot allocation. The server allocates an active slot at time t to user i such that

$$i = \operatorname{argmax}_{j: q_{j,t}=1, b_{j,t} > 0} c_j p_j,$$

i.e., among the non-empty and connected queues the policy selects the queue with the largest index $c_i p_i$, which may be viewed as the immediate expected cost reduction (In case of multiple slot assignment we assume that the connectivity and backlog for each slot are known. More will be discussed in Section III-B).

This greedy policy is in general not optimal. An example can be found in [18]. In the next section we provide sufficient conditions under which this greedy policy is optimal under different assumptions on the user connectivity processes.

B. Summary of Notations and Assumptions

We consider time evolution in discrete time steps indexed by $t = 0, 1, \dots, T - 1$, with each increment representing a time slot length. Slot t refers to the time interval $[t, t + 1)$. In subsequent discussions we will use terms *slots*, *steps* and *stages* interchangeably. A frame consists of an active interval followed by an inactive interval. The first frame starts from $t = 0$ and ends at $t = LM$ which is the LM -th time slot. The second frame starts at $t = LM$ and so on.

We will use subscripts to denote the time index and to denote a specific user/queue. For example $b_{i,t}$ denotes the buffer occupancy at the beginning of time slot t for the i -th queue. All boldface letters represent column vectors and all normal letters represent scalars/random variables. Whenever we need to distinguish two policies, we show the policy as a superscript. For example $b_{i,t}^\pi$ means the backlog of queue i at time t under policy π .

A list of important notations is as follows.

M : The length of the active period in number of slots.

L : the length of a cycle in multiples of M slots, i.e., a cycle has a length of LM slots. Equivalently, L is the inverse of the duty cycle.

$\mathbf{b}_t = [b_{1,t}, b_{2,t}, \dots, b_{N,t}]'$: The column vector of all queue occupancies at time t .

$\mathbf{b}_t^{i+} = \mathbf{b}_t + \mathbf{e}_i$, where \mathbf{e}_i is the N dimensional vector with all the values being zero except a one in the i -th position.

$\mathbf{a}_t = [a_{1,t}, a_{2,t}, \dots, a_{N,t}]'$: The number of packet arrivals during time slot t .

q_t : The channel connectivity during the t -th time slot.

p_i : The transmission success probability of queue i .

\mathcal{F}_t : The σ -field of the information available up to time t .

Below we summarize important assumptions underlying our network model.

- 1) We assume that each user has an infinite buffer. Without this assumption we need to introduce penalty for packet dropping/blocking. This is an important extension to the work presented here but is out of the scope of this paper.
- 2) We assume that the slot allocation for active period starting at time t cannot be used to transmit the possible packet arrival during the t^{th} slot, i.e., within $[t, t + 1)$. This is because the exact arrival time of this packet is random, and unless it arrives right before t it cannot be transmitted during that slot.
- 3) We assume that the channel state does not change during an active time slot.
- 4) We assume that the acknowledgments are immediate (i.e. we find out whether a transmission is successful or not at the beginning of the next slot.)

III. SUFFICIENT CONDITIONS FOR THE OPTIMALITY OF THE GREEDY POLICY

In this section we study the optimality of the greedy policy discussed earlier. To make the discussion simpler we start by considering the case where an active period consists of a single slot $M = 1$. We then generalize the results to the case where $M > 1$. We also assume that T is an integer multiple of LM , the length of a cycle. This assumption allows us to keep the results in a simple form, but it can be easily relaxed.

Due to space limit the proofs of the lemmas and theorems are not presented, but they can be found in [18].

A. Single Slot Active Period

Let $M = 1$, i.e. there is only one slot to allocate during an active period. Note that in this case active time slots are those at $t = 0, L, 2L, \dots, \frac{T}{L} - 1$. The following lemma holds true regardless of any assumptions on the connectivity of the queues.

Lemma 1: Let π be the optimal policy for state \mathbf{b}_0 and let π' be the optimal policy for state \mathbf{b}_0^{i+} . Then we have,

$$E^{\pi'}[C|\mathcal{F}_0, \mathbf{b}_0^{i+}] - E^\pi[C|\mathcal{F}_0, \mathbf{b}_0] \leq \frac{c_i(1 - \beta^T)}{1 - \beta}. \quad (2)$$

Below we consider two models for the channel connectivity and derive lower bounds on the value $E^{\pi'}[C|\mathbf{b}_0^{i+}] - E^\pi[C|\mathbf{b}_0]$.

1) *No Information on Connectivity*: In this part we assume the following about the channel connectivity.

No-info - At the beginning of each active slot, the server is informed about the connectivity for that slot, but the server does not know the statistics of the connectivity process, e.g., it does not know how the connectivity changes from one time slot to the other.

Lemma 2: Let π be the optimal policy for the initial state \mathbf{b}_0 and let π' be the optimal policy for state \mathbf{b}_0^{i+} . If there is no information about the channel connectivity process, then we have

$$\begin{aligned} & E^{\pi'}[C|\mathcal{F}_0, \mathbf{b}_0^{i+}] - E^\pi[C|\mathcal{F}_0, \mathbf{b}_0] \\ & \geq \frac{r_i(1-p_i)(1-(\beta^L(1-p_i))^{\frac{T}{L}})}{1-\beta^L(1-p_i)}, \end{aligned} \quad (3)$$

where $r_i = c_i + \beta c_i + \dots + \beta^{L-1} c_i$.

Theorem 1: Suppose the initial backlog state is \mathbf{b}_0 and suppose queues i and j are connected and non-empty. Let π be the policy that allocates the slot to queue i and let π' be the policy that allocates the slot to queue j . If there is no information available on the statistics of channel connectivity (**No-info**), but only that they are both connected in the current slot, then we have

$$E^\pi[C|\mathcal{F}_0, \mathbf{b}_0] \leq E^{\pi'}[C|\mathcal{F}_0, \mathbf{b}_0],$$

if

$$\begin{aligned} & p_i r_i + \beta^L p_i \left(\frac{r_i(1-p_i)(1-(\beta^L(1-p_i))^{\frac{T}{L}-1})}{1-\beta^L(1-p_i)} \right) \\ & \geq p_j r_j + \beta^L p_j \frac{c_j(1-\beta^{T-L})}{1-\beta}. \end{aligned} \quad (4)$$

(Note that the right hand side is simply equal to $\frac{p_j c_j (1-\beta^T)}{1-\beta}$, but we leave it in this form to make it easier to compare the two sides).

Corollary 1: Suppose the state at $t = 0$ is \mathbf{b}_0 and suppose queue i is connected and non-empty. If there is no information on the statistics of the channel connectivity process, then it is optimal to allocate the slot at $t = 0$ to queue i if (4) holds for all $j \neq i$ such that $q_{j,0} = 1$ and $b_{j,0} > 0$.

2) *Independent Connectivity*: In this part we assume the following about the channel connectivity.

Indep - At each active time slot, user i is connected to the server with probability q_i independent of all past history. The quantities q_i are known to the server. In addition, at the beginning of each active time slot the server knows whether a queue is connected for that slot.

This assumption is valid if for example the length of the inactive period is very large in comparison with the channel variations, so that the channel states during successive active periods appear independent.

Lemma 3: Let π be the optimal policy for state \mathbf{b}_0 and let π' be the optimal policy for state \mathbf{b}_0^{i+} . If the channel changes state independently at the beginning of each active

slot (**Indep** -), then we have

$$\begin{aligned} & E^{\pi'}[C|\mathcal{F}_0, \mathbf{b}_0^{i+}] - E^\pi[C|\mathcal{F}_0, \mathbf{b}_0] \\ & \geq \frac{r_i(1-p_i q_i)(1-(\beta^L(1-p_i q_i))^{\frac{T}{L}})}{1-\beta^L(1-p_i q_i)}, \end{aligned} \quad (5)$$

where $r_i = c_i + \beta c_i + \dots + \beta^{L-1} c_i$.

Theorem 2: Suppose the initial state is \mathbf{b}_0 and suppose queues i and j are connected and non-empty. Let π be the policy that allocates the slot to queue i and let π' be the policy that allocates the slot to queue j . Using the channel model defined by **Indep**, we have

$$E^\pi[C|\mathcal{F}_0, \mathbf{b}_0] \leq E^{\pi'}[C|\mathcal{F}_0, \mathbf{b}_0],$$

if the following inequality holds:

$$\begin{aligned} & p_i r_i + \beta^L p_i \left(\frac{r_i(1-p_i q_i)(1-(\beta^L(1-p_i q_i))^{\frac{T}{L}-1})}{1-\beta^L(1-p_i q_i)} \right) \\ & \geq p_j r_j + \beta^L p_j \frac{c_j(1-\beta^{T-L})}{1-\beta}. \end{aligned} \quad (6)$$

Corollary 2: Suppose the state at $t = 0$ is \mathbf{b}_0 and suppose queue i is connected and non-empty. If the channel model is as **Indep**, then it is optimal to allocate the slot at $t = 0$ to queue i if (6) holds for all $j \neq i$ such that $q_{j,0} = 1$ and $b_{j,0} > 0$.

Remark 1: Note that the sufficient condition in Theorem 2 (Equation (6)) is weaker than the one in Theorem 1 (Eqn (4)), i.e. it is satisfied more easily. Essentially the information on the connectivity process allows us to derive a tighter bound for the optimality of the greedy policy.

Remark 2: As L increases the sufficient conditions (4) and (6) become weaker. Specifically in the limit as $L \rightarrow \infty$ it can be seen that it is optimal to serve queue i if $p_i c_i \geq p_j c_j$ for all $j \neq i$ which is essentially the greedy policy. Therefore in this case the sufficient conditions are tight and the greedy policy is optimal.

Remark 3: Although all theorems in this section are based on the optimal allocation at time $t = 0$, it can be seen that all the results can be easily extended to the bandwidth allocation at time t by replacing T with $T - t$ in all the sufficient conditions. This is due to the fact that T is essentially the ‘‘time to go’’ in all these results and if we start at time t , then the time to go is $T - t$.

B. Multiple Slot Active Period

In this part we assume $M > 1$ and find the sufficient conditions for the optimality of the greedy policy, using the same channel connectivity models defined in Section III-A. Note that in the case of $M > 1$, the allocation decision made by the server is delayed, in the sense that the server uses the backlog information at time t to make the allocation decision for time $t, t+1, \dots, t+M-1$. By the time the m -th slot ($m > 1$) is used (at time slot $t+m-1$), the backlog of the queues may have changed. In order to avoid complications caused by this information delay we make the following assumption.

Assumption 1: When $M > 1$ the server makes the allocation decision for each active time slot individually. At each active time slot the server knows the backlog and connectivity of all the queues during that time slot before making the allocation decision.

This assumption certainly holds in the case of down-link communication (from the server to the users). In the uplink, as long as the length of the active period is small compared to the arrival probability, this is a good approximation. In any case this assumption introduces a lower bound on the cost of the real system (with information delay). Note that Lemma 1 holds in the case of multiple slot active period as well.

Lemma 4: Let π be the optimal policy for state \mathbf{b}_0 and let π' be the optimal policy for state \mathbf{b}_0^{i+} . If the channel state process is not known, then we have

$$\begin{aligned} & E^\pi [C|\mathcal{F}_0, \mathbf{b}_0^{i+}] - E^{\pi'} [C|\mathcal{F}_0, \mathbf{b}_0] \\ & \geq \frac{r'_{i,1}(1 - (\beta^{LM}(1 - p_i)^M)^{\frac{T}{LM}})}{1 - \beta^{LM}(1 - p_i)^M}, \end{aligned}$$

where

$$r'_{i,m} = \sum_{k=m-1}^{M-1} \beta^k (1 - p_i)^{k-m+2} c_i + \sum_{k=M}^{LM-1} \beta^k (1 - p_i)^M c_i.$$

Theorem 3: Let t be the m -th active slot of an active period and let $t' = t - m + 1$ (this is the first slot of the active period). Suppose the state at t is \mathbf{b}_t and suppose queue i is connected and non-empty. If the channel model is as defined by **No-info** and there are M slots per active period, then it is optimal to allocate the slot at time t to queue i if the following inequality holds for all $j \neq i$ such that $q_{j,t} = 1$ and $b_{j,t} > 0$:

$$\begin{aligned} & \frac{p_i r'_{i,m}}{1 - p_i} + \beta^{LM} p_i \frac{r'_{i,1}(1 - (\beta^{LM}(1 - p_i)^M)^{\frac{T-t'+1}{LM}-1})}{1 - \beta^{LM}(1 - p_i)^M} \\ & \geq \frac{p_j r'_{j,m}}{1 - p_j} + \beta^{LM} p_j \frac{c_j(1 - \beta^{T-t'-LM})}{1 - \beta}. \end{aligned} \quad (7)$$

The following results are for the case of independent channel connectivity (**Indep**).

Lemma 5: Let π be the optimal policy for state \mathbf{b}_0 and let π' be the optimal policy for state \mathbf{b}_0^{i+} . If the channel changes state independently at the beginning of each active slot, then we have

$$\begin{aligned} & E^\pi [C|\mathbf{b}_0^{i+}] - E^{\pi'} [C|\mathbf{b}_0] \\ & \geq \frac{r''_{i,1}(1 - (\beta^{LM}(1 - p_i q_i)^M)^{\frac{T}{LM}})}{1 - \beta^{LM}(1 - p_i q_i)^M}, \end{aligned}$$

where

$$\begin{aligned} r''_{i,m} & = \sum_{k=m-1}^{M-1} \beta^k (1 - p_i q_i)^{k-m+2} c_i \\ & \quad + \sum_{k=M}^{LM-1} \beta^k (1 - p_i q_i)^M c_i. \end{aligned}$$

Theorem 4: Let t be the m -th active slot of an active period. Let $t' = t - m + 1$. Suppose the state at t is \mathbf{b}_t and suppose queue i is connected and non-empty. If the channel model is as defined by **Indep** and there are M slots per active period, then it is optimal to allocate the slot at $t = 0$ to queue i if the following inequality holds for all $j \neq i$ such that $q_{j,t} = 1$ and $b_{j,t} > 0$:

$$\begin{aligned} & p_i \frac{r''_{i,m}}{1 - p_i q_i} \\ & + \beta^{LM} p_i \frac{(r''_{i,1}(1 - (\beta^{LM}(1 - p_i q_i)^M)^{\frac{T-t'}{LM}-1})}{1 - \beta^{LM}(1 - p_i q_i)^M)} \\ & \geq \frac{p_j r''_{j,m}}{1 - p_j} + \beta^L p_j \frac{c_j(1 - \beta^{T-t'-LM})}{1 - \beta} \end{aligned} \quad (8)$$

Remark 4: All the results presented in this section hold for all values of T . Specifically one can let $T \rightarrow \infty$ to derive the sufficient conditions for the optimality of the greedy policy in the case of an infinite horizon.

IV. NUMERICAL ANALYSIS

As shown earlier, the greedy policy is not necessarily optimal and in the previous section we found sufficient conditions for its optimality. In this section we employ this policy, regardless of whether it is optimal for the scenarios considered, and study the performance of this policy via a few numerical examples. In particular, we are interested in (1) how the performance of this policy (in terms of packet holding cost) varies as the duty cycle ($\frac{1}{L}$) changes while fixing the active period M ; and (2) how the performance varies as M changes while fixing L (i.e., fixing the duty cycle but varying the frequency of cycling).

Note that as L increases the duty cycle decreases and therefore we expect the total cost to increase. On the other hand large L means longer inactive intervals, which implies longer lifetime of the system. Therefore it is important to see how the performance degrades as the lifetime increases. The effect of M is more complicated. As M increases (for fixed L) the system has longer cycles, i.e., switches between on and off periods less often. This in turn increases the system lifetime as turning devices on and off typically consumes nonnegligible energy especially for low power devices. But at the same time, longer cycles may increase the probability that an active slot coincides with empty queues which causes performance degradation.

We assume that the channel states in different slots are independent, with a fixed connectivity probability. The server does not need to know this probability (in fact the greedy policy does not require any information about how the state changes, it only needs to know the current state and the current backlog). While it is obvious that the example considered in this section is certainly not sufficient for a full characterization of the behavior of the system in general, it nevertheless provides some interesting insight on the properties of the greedy policy and the effect of parameters like M and L .




Fig. 2. Cost Variation as a function of L and M

Consider the following scenario. There are two queues. Arrivals are Bernoulli, i.e. at each time slot there is an arrival to queue i with probability u_i and there are no arrivals with probability $1 - u_i$. Other parameters are as follows: $\beta = 0.999$, $T = 1000$, $c_1 = 15$, $c_2 = 5$, $u_1 = 0.05$, $u_2 = 0.05$, $p_1 = 0.9$, $p_2 = 0.6$, $q_1 = 0.7$, and $q_2 = 0.9$. Figure 2-a illustrates the total cost as a function of L . Different curves correspond to different values of M . The curve is an average over 200 simulations.

Note that in this case as M increases the curve shifts up. This clearly shows that small M , i.e., shorter cycles or more frequent on/off switching performs better. This is not always the case. In [18] we provide another example where queue 1 has a smaller success probability and a smaller connectivity probability and observe that increasing M has a much smaller effect on the cost performance and the frequency of cycling made little difference in terms of costs for the same duty cycle. This is because in this case more packets are queued up in the system and the queues do not become empty as often as the previous scenario. Consequently, multiple slot allocation at a time ($M > 1$) is well utilized in that there is a high probability that these slots are used and the queues are less likely to become empty (see [18] for more discussion).

V. CONCLUSION

In this paper we analyzed the optimality of an index/greedy policy for allocating time slots in a low duty cycled system. Each user is associated with a connectivity variable and a transmission success probability. The greedy policy allocates the channel to the non-empty connected queue with the largest immediate expected cost. We provided sufficient conditions for this greedy policy to be optimal. We then studied the effect of changing the duty cycle and the number of active slots per frame on the performance of the greedy policy. The performance degrades as the duty cycle decreases. When the system is near the boundary (the queues are empty most of the time) increasing the number of active slots per frame (M) degrades the performance. On the other hand decreasing the duty cycle and increasing

M increases the lifetime of the system. Comprehensive modeling of this tradeoff is part of our future study.

REFERENCES

- [1] J. S. Baras, A. J. Dorsey, and A. M. Makowski, "Two competing queues with linear costs and geometric service requirements: The μc rule is often optimal," *Adv. Appl. Prob.*, vol. 17, pp. 186–209, 1985.
- [2] C. Buyukkoc, P. Varaiya, and J. Warland, "The $c\mu$ -rule revisited," *Advances in Applied Probability*, vol. 17, pp. 237–238, 1985.
- [3] J. G. Shanthikumar and D. D. Yao, "Multi class queueing systems: Polymatroid structure and optimal scheduling control," *Oper. Res.*, vol. 40, pp. 293 – 299, 1992.
- [4] J. Edmonds, "Submodular functions, matroids and certain polyhedra," *In Proc. of the Calgary International Conference on Combinatorial Structures and Their Applications*, Gordon and Breach, New York, pp. 69–87, 1970.
- [5] D. Bertsimas and J. Nino-Mora, "Conversion laws, extended polymatroids and multi-armed bandit problems," *Mathematics of Operations Research*, vol. 21, pp. 257–306, 1996.
- [6] T. C. Green and S. Stidham Jr., "Sample path conservation laws, with applications to scheduling queues and fluid systems," *Queueing Systems*, vol. 36, pp. 175–199, 2000.
- [7] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 466–478, March 1993.
- [8] L. Tassiulas, "Scheduling and performance limits of networks with constantly changing topology," *IEEE Transactions on Information Theory*, vol. 43, no. 3, pp. 1067–73, May 1997.
- [9] N. Bambos and G. Michailidis, "On the stationary dynamics of parallel queues with random server connectivities," *Proc. 43th Conference on Decision and Control (CDC)*, pp. 3638–43, 1995, New Orleans, LA.
- [10] C. Lott and D. Teneketzis, "On the optimality of an index rule in multi-channel allocation for single-hop mobile networks with multiple service classes," *Probability in the Engineering and Information Sciences*, vol. 14, no. 3, pp. 259–297, July 2000.
- [11] M. J. Neely, E. Modiano, and C. E. Rohrs, "Power allocation and routing in a multibeam satellite with time-varying channels," *IEEE Proceedings of INFOCOM*, 2002.
- [12] M. J. Neely, E. Modiano, and C. E. Rohrs, "Dynamic power allocation and routing for time-varying wireless networks," *IEEE/ACM Transactions on Networking*, Vol. 11, No. 1, pp. 138–152, 2003.
- [13] B. Hajek, "Optimal control of two interacting service stations," *IEEE Trans. Auto. Control*, AC-29, pp. 491–499, 1984.
- [14] P. Whittle, "Restless bandits: Activity allocation in a changing world," *A Celebration of Applied Probability*, ed. J. Gani, *Journal of applied probability*, vol. 25A, pp. 287–298, 1988.
- [15] R. Weber and G. Weiss, "On an index policy for restless bandits," *Journal of Applied Probability*, vol. 27, pp. 637–648, 1990.
- [16] J. Nino-Mora, "Restless bandits, patial conservation laws, and indexability," *Advances in Applied Probability*, Vol. 33, no. 1, pp. 76–98, 2001.
- [17] C. H. Papadimitriou and J. N. Tsitsiklis, "The complexity of optimal queueing network control," *Mathematics of Operations Research*, Vol. 24, No. 2, pp. 293–305, May 1999.
- [18] N. Ehsan and M. Liu, "Dynamic bandwidth allocation for low power devices with random connectivity," *EECS Technical Report CSPL 369*, University of Michigan, Ann Arbor, 2005.