Energy-Efficient Communication in Battery-Constrained Portable Devices

Athina Markopoulou, Yan Li, Carri Chan, Nick Bambos

Electrical Engineering Department
Stanford University, CA, USA
{amarko,liyan, cwchan, bambos}@stanford.edu

***Abstract*— Portable devices (such as personal digital assistants and laptops with wireless connectivity) are becoming ubiquitous. As their functionality and capabilities increase, their energy consumption requirements also increase. Yet, these devices have to operate on limited batteries. In order to maximize the battery lifetime, it is necessary to optimize the use of energy at various components of such a device. In this paper, we consider a single portable device operating on a limited battery that transmits information over an interference-limited wireless channel. We seek to optimize the power consumption on the communication radio in this device, by controlling both the operation mode and the transmission power. We model the general problem using dynamic programming, obtain the optimal solutions for insightful special cases and explore various design tradeoffs. Our work provides an analytical framework for stochastic modeling and optimization of energy spent for communications in battery-operated portable devices.**

I. Introduction

Recent advances in the design of portable devices combined with advances in wireless access and network convergence, have made portable - and in particular handheld- devices extremely popular. Such devices conveniently provide many services including communication, computation, personal information management, and Internet access. For example, laptops equipped with Wi-Fi capability provide today ubiquitous access to the Internet; traditional cellphones are enhanced with Internet-access and other services; personal digital assistants (PDAs) combine various built-in capabilities (including cellular, WiFi and VoIP, storage and computation) and gain more and more momentum. We are particularly interested in this last category, which has the potential to become the dominant integrated solution in this space.

Portable devices operate on a limited battery. Clearly, it is desirable to maximize the battery lifetime for the user's convenience (e.g. to maximize the time to recharging). However, limitations in the size and weight of portable (and particularly of handheld) devices result in limitations in the allowed battery size. Therefore, it is important to efficiently use the battery reserves so as to achieve high-performance operation and long battery lifetime. This optimization complements improvements in the size-efficiency of batteries.

Power management and control can be applied on various components and layers of a battery-constrained portable device. In this paper, we focus on the efficient use of the communication radio, and we control both (i) the transitions between the various operation modes (e.g. *on*, *sleep*, *off*) and (ii) the power used for transmission. In this first step, our work develops the analytical framework for stochastic modeling and optimization of energy efficient communication, taking into account both control aspects. We formulate the general problem in a dynamic programming framework, and explore various design tradeoffs through analysis of insightful special cases. We hope that this approach will be used by the research community to model additional cases, explore the design tradeoffs, and develop and evaluate practical heuristics.

The rest of the paper is structured as follows. In section II we discuss related work in the area of energy-efficient communications, and where our work lies in this problem space. In section III, we introduce the general model for a single node and discuss the various design tradeoffs. In section IV, we discuss variations of the model to address several operational issues and special case scenarios. In sections V and VI, we present and analyze two insightful special cases, which, apart from serving as a case study, also sheds light to some fundamental tradeoffs of the general problem. In particular, in section V, we study the case where the node is always on and we control the transmission power; we also optimally choose the initial power reserves. In section VI, we control both the transmission power and the operation mode of the radio. Section VII compares the policies obtained for the two special cases and for benchmark systems. Section VIII concludes the paper.

II. Related Work

There has been a large body of work related to energy-efficiency in various contexts, including system-level power management, wireless ad-hoc and sensor networks.

System-level *dynamic power management (DPM)* algorithms, preserve energy by switching idle components to lower power consumption states, e.g. see [1]; clearly there is a tradeoff between energy savings and performance. In [2], this problem has been studied as a stochastic optimization problem; system resources are modeled by states, capturing the power-performance tradeoff, and transitions between states can be optimally managed by a policy manager. In the past, timeout-based heuristics, [3], (originating from earlier work on hard disks and interactive terminals) as well as predictive techniques, [4], have also been extensively used.

On the other hand, in the context of wireless communications with interference-limited environments, there has been an extensive amount of work on *transmission power control (TPC)*; e.g. for an example of power control for packet-based traffic see [5]. In that context, the purpose is to optimally control the transmission power level to combat interference, so as to use low power while still achieving a desired quality-of-service level. In addition to saving energy, TPC has additional benefits, such as increase in network capacity (by keeping interference and thus stress to the channel low) and decrease of exposure to electromagnetic radiation.

Most past and ongoing work on energy-efficiency for handheld and wireless 802.11 devices, focuses on the DPM aspect, that is on when to switch the radio between various operation modes; e.g. see [6], [7] for a protocols and systems perspective and [8] for a stochastic optimization treatment of the problem. On the contrary, little attention is given on the transmission power aspect, as e.g. in [9], [10]. The reason is that with the current transceivers, the energy spent while in high consumption states (e.g. *on*) is significantly higher than in the low consumption ones *off*, *sleep*. This is mainly due to the design of power amplifiers, whose efficiency is non-linear with the transmission power: they are efficient for high transmission power and inefficient for low transmission power; a discussion in the context of 802.11 [11] can be found in [10]. However, as power control starts getting implemented, e.g. in 802.11h [12], power amplifiers are improving and becoming more efficient even in lower power levels. For example, in [9], [10], TPC is used in combination with PHY rate adaptation for 802.11a/h.

In this paper, we aim at minimizing the energy spent for communication by the radio of a single battery-constrained device; we seek to jointly control both the operation state (which resembles DPM) and the transmission power (which resembles TPC). The relative benefit from DPM vs. TPC eventually depends on the scenario itself; e.g. as power amplifiers become more efficient DPM and TPC become comparable; for long range transmission (due to geographical or military constraints) the transmission power dominates; finally, TPC brings additional benefits other than battery, such as the increase in network capacity.

Energy efficiency is also critical in the context of sensor networks, for which resources are even more limited. A nice survey on energy efficient techniques for sensor networks is given in [13]. The characteristics of radios for sensor networks differ from traditional radios (e.g. in that *Tx* power is comparable to *Rx*, *Idle* power). Because the dominant part of power is spent for keeping the radio on, the widely used technique for saving energy, is again to turn off the radio during idle periods.

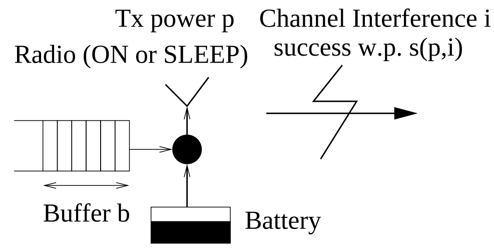

Fig. 1. A single portable device with battery reserves π and b packets in the buffer. The wireless channel has interference i. The radio has control over (i) its operation mode (ON, SLEEP, etc) and over (ii) the transmission power p that is used to transmit a packet.

Finally, another body of work in wireless ad-hoc and sensor networks, takes a network-centric point of view and tries to maximize the lifetime not only of individual nodes but also of the network as a whole. Issues include how to establish topology, how to route traffic under stringent energy constraints, and how to manage the operational state transitions while maintaining undisturbed operation. Representative examples of work in this area, include - but are not limited to - [14], [15], [16], [17]. In [18], closed-loop control concepts were used for power management of networks-on-chips.

Methodologically, our work resembles more the stochastic modeling and optimization approaches, e.g. [2], [8]. However, we control both the radio state (DPM) and the transmission power (TPC); in the following sections, we explore how the optimal control depends on the relative values of operating power (in *on*, *sleep* states) vs. transmission power.

III. GENERAL MODEL AND PROBLEM FORMULATION

In this section we introduce the basic model, reflecting our problem formulation and capturing the performance tradeoffs and control issues. We embed the problem within a Markov decision process framework and use dynamic programming to compute the optimal control [19].

A. System Description

We consider the single portable device shown in Fig. 1. The device has a radio which can be in one of two modes: either ON or SLEEP. (This can be extended to a larger number of states, at the cost of additional computation complexity, as discussed later. Without loss of generality, let us consider two states for the moment). The device also has a battery with energy reserves π, and a buffer with b packets containing data that the user has accumulated and wants to transmit over the wireless channel (e.g. to send emails or any other data transfer). The radio can transmit packets only when it is ON. The wireless channel has interference i, which affects the probability of successful packet transmissions. The device wants to transmit all packets across the wireless channel using the minimum amount of battery energy. This can be achieved by appropriately controlling the operating mode of the radio and its transmission power p. Time is slotted and indexed by $t = 0, 1, 2, 3, ...$.

Using transmission power p, a packet is successfully transmitted with probability $s(p, i)$ that depends on the interference i (or quality of the channel). Clearly, an unsuccessful transmission wastes battery energy. The interference i in the channel fluctuates according to a time-homogeneous Markov chain, taking values in the finite set I of all attainable interference states. It switches with probability q_{ij} from state $i \in I$ in a time slot to state $j \in I$ in the next time slot. Furthermore, it is assumed that the interference is not responsive to transmitter actions.

Initially the device has battery π_0 and b_0 packets. We assume that the packets are already collected and stored in the buffer, and the node wants to transmit them over the wireless channel, e.g. to its neighbors or to the base station. In general, the user of the device may continuously produce new data packets that arrive to its buffer. The user activity, which can be thought as arrival process to the buffer, follows its own duty-cycle, which might also be subject to another optimization. For the moment, we focus on how to manage the radio in order to transmit a certain amount of already stored packets, while spending the minimum amount of energy. Later on, we can modify our model to include a live source or continuous activity.

The dilemma faced by the radio is the following. On one hand it wants to transmit the packets as soon as possible; in order to do so though, it has to be ON, spending operating power P_{on}, and a certain transmission power (p). On the other hand, it wants to avoid spending any power in order to preserve its limited battery. Clearly, there is a tradeoff between these two conflicting goals. We seek the optimal power management to jointly optimize these goals.

B. Discussion of Costs

There are several pressures to be considered and captured into performance/operational costs.

Packet Delivery and Costs. We assume that b_0 packets, containing information produced by the user and initially stored in the buffer. In order to model our intention to transmit these data on time, we introduce the following costs. The backlog cost $B(b)$, incurred at each time slot, models the urgency to deliver the packets as soon as possible. In addition, if the system terminates (because it runs out of battery) without delivering all data, we introduce a terminal cost (b^x) associated with the number of remaining packets.

Power Control and Costs. In order for the device to transmit the packets, it needs to spend energy from its battery reserves, which initially is π_0. The following power[1] costs are associated with the operation of the device:

- *Power Spent on Operation Modes.* The radio spends operating power P_{on} and P_{sleep}, for every time slot in mode ON and SLEEP respectively. In general, P_{on} is the dominant part, while $P_{sleep} \simeq 0$ is negligible; e.g. the ratio P_{on}:P_{sleep} depends on the specific device and radio (typically in the order of 10:1 in 802.11 radios). Additional operational modes of interest can easily be incorporated in this model by increasing the state space. For example, 802.11b radios have the *awake, sleep, off* modes and the *awake* state can further be divided into *transmit, receive*(which is in general non negligible) and *idle*. For simplicity, in this paper, we refer only to *on* and *sleep* modes. Methodologically, the same analysis can be applied to any number of states, at the additional cost of increased computational complexity.
- *Power Spent on Transitions between Modes.* In general, switching between modes requires to spend transmission power. Typically, more power is required to wake-up the node. [2]
- *Transmission Power.* The radio can use transmission (Tx) power p from a bounded range of $p = 0, 1, 2, ...P_{max} \leq \pi - P_{on}$. The specific values depend on the standard, e.g. see [12] for the possible power levels in 802.11h.
- *Stress Induced to the Channel.* Transmission power p when the interference is i, introduces a cost $\Psi(p, i)$ paid in that slot. This cost may reflect the interference stress that the transmission under consideration induces on the channel, e.g., interfering with 'background' transmissions from other transmitters that use the same channel. The latter may in turn stress the original 'foreground' transmitter in response to its power increases, thus generating more interference on it. This entanglement effect is implicitly captured in the cost $\Psi(p, i)$. The cost $\Psi(p, i)$ should be increasing in both p and i, consistently with the intuition that the more congested the channel is, the more power should be spent to capture it and support the required success probability. The stress that transmitting power induces to the channel may be critical in densely populated wireless networks with a large number of transmitters. Recent work characterized the capacity of such interference - limited wireless ad-hoc networks [20]. A final consideration is how Power Control interacts with Multiple Access protocols such as CSMA/CA; this problem is important, but outside the scope of this paper.
- *Cost of Unused Battery.* All the above-mentioned power requirements drain the battery reserves π_0. We would like to transmit the packets, using the minimum re-

[1]Throughout the text, we will sometimes use the terms "power" and "energy" interchangeably, considering a unit time of 1.

[2]We further assume that transitions happen with negligible transition delay. This is the case when the transition time is shorter than the duration of the time slot. One could also model a transition delay that lasts for several time slots, by maintaining additional the state to keeping track in a time window, from the time a transition is requested until the transition takes place. We omit it here for simplicity.

quired battery. If initially we “payed” for battery π_0, but in practice we sent all packets using only $\pi < \pi_0$, then we associate a terminal cost π^y with the unused battery, to express our intention to make an initial investment π_0 that is just sufficient for the job, no more no less. This investment can be thought in terms of money, engineering effort to manufacture a battery larger than needed (given size and form constraints and desirable lifetime for the user's convenience), or engineering effort to scavenge energy from the environment. The more general question is *how much energy we need per packet,* $\lim_{b_0 \to \infty} \pi_0 / b_0$*?*

Depending on the relative values of p and P_{on}, most of the energy savings can come from DPM or TPC; clearly the larger $P_{on} : p$ the more the benefit from DPM. However, even when $p << P_{on}$ (as for example in the motes used in sensor networks), the stress $\Psi(p, i)$ induced to the channel may be comparable to P_{on}. Furthermore, transmission power control is implicitly entangled with DPM control because it affects how long we keep the radio ON and pay the P_{on} operating cost. Our framework enables us to explore tradeoffs like P_{on} vs. p.

C. System State and Optimal Control

The objective is to transfer all packets stored in the buffer, minimizing the overall cost. The system state to be tracked in each time slot is:

$$(b, \pi, i, mode)$$

that is, the current number b of packets left in the transmitter buffer, the remaining battery π, the current interference state i in the channel and the $mode$ of the radio (ON or SLEEP) in the previous time slot. The controls applied, or decisions made, in each time slot, are (m, p), where:

- m is the mode of the system in the current time slot: $m \in \{ON, SLEEP\}$.
- p is the transmission power, and can be chosen from a bounded range $[0, P_{max}]$, provided that the current mode is ON and there is enough battery left ($p \leq \pi - P_{on}$).

System Evolution. The device initially starts with b_0 packets in the buffer. In every time slot, the optimal control chooses the mode in the current time slot m and the transmission power p. The evolution of the system ends when either all packets are transmitted ($b = 0$), or the battery is emptied ($\pi = 0$). Given this formulation, the system simply becomes a controlled Markov chain. Hence, we can develop a Dynamic Programming (DP) recursion to compute the optimal control [19].

Let $J(b, \pi, i, mode)$ be the cost-to-go, that is the minimum cost incurred from now on until termination, given that the optimal control (m, p) is used and the current state is $(b, \pi, i, mode)$. The quantity $J(b, \pi, i, mode)$ satisfies the following functional recursive equations, (1) and (2), for $b \in \{0, 1, , ..., b_0\}$,$\pi$, $i \in I$, $mode \in \{on, sleep\}$.

$$\begin{aligned} J(b, \pi, i, sleep) = \min_m \{ & B(b) + \\ & 1_{m=sleep} \sum_{j \in I} q_{ij} J(b, \pi, j, sleep) + \\ & 1_{m=on} [P_{sleep \to on} + \sum_{j \in I} q_{ij} J(b, \pi - P_{sleep \to on}, j, on)] \} \end{aligned} \tag{1}$$

Eq. 1 describes the evolution, given that we are in SLEEP mode. In the current time slot, we are paying the backlog cost. If we decide to stay in SLEEP mode, [3] the only state change is that interference goes from i to j. Alternatively, we may decide to switch back ON, paying transition (waking up) cost $P_{sleep \to on}$ in the current time slot and future expected cost $\sum_{j \in I} q_{ij} J(b, \pi - P_{sleep \to on}, j, on))$.

$$\begin{aligned} J(b, \pi, i, on) = \min_{m,\ p} \{ & B(b) \\ & + 1_{m=sleep} \{P_{on \to sleep} + \sum_{j \in I} q_{ij} J(b, \pi - P_{on}, j, sleep)\} \\ & + 1_{m=on} \{P_{on} + p + \Psi(p, i) \\ & + s(p, i) \sum_{j \in I} q_{ij} J(b - 1, \pi - p - P_{on}, j, on) \\ & + [1 - s(p, i)] \sum_{j \in I} q_{ij} J(b, \pi - p - P_{on}, j, on)\}\} \end{aligned} \tag{2}$$

Eq. 2 corresponds to the case that we are in the ON mode. Backlog cost $B(b)$ is payed at the current time slot. Additional cost will be payed, depending on the decision to stay ON or switch to SLEEP. The optimal control will choose the transition that minimizes the total expected (current and future) cost.

The 2^{nd} line of Eq. 2 corresponds to the case that we decide to switch from ON to SLEEP. In the current time slot, we pay backlog and operating cost. The future expected cost depends on the channel transitions (from i to j) and on the cost-to-go in the SLEEP state, given in Eq. 1: $\sum_{j \in I} q_{ij} J(b, \pi - P_{on}, j, sleep)$.

The 3^{rd} and 4^{th} lines of Eq.2 correspond to the case that we decide to stay ON. Then the optimal cost-to-go cost, exercising optimal control, is comprised of the cost payed at the current slot (backlog cost $B(b)$ and power cost $P_{on} + p + \Psi(p, i)$) and two more terms depending on the outcome of the current transmission. $J(b - 1, \pi - p - P_{on}, j, on)$ corresponds to the case that the transmission is successful; the number of packets is reduced by one and the battery is reduced by the operating cost P_{on} and the transmission cost p. $J(b, \pi - p - P_{on}, j, on)$ corresponds to the case that the transmission is unsuccessful, thus the buffer level stays the same, although we spent the same amount of battery $p + P_{on}$, as before.

[3]Notation: the indicator function $1_{m=on}$ takes the value 1 if $m = on$ and 0 otherwise. Similarly, $1_{m=sleep} = 1$ if $m = sleep$ and 0 otherwise.

Terminal Costs. The recursion terminates either when we run out of battery ($\pi = 0$) or when we transmit all packets ($b = 0$). The terminal costs depend on the remaining packets and battery respectively:

$$J(b, 0, i, mode) = b^x \tag{3}$$

$$J(0, \pi, i, mode) = \pi^y \tag{4}$$

The parameters x, y capture how much we value the importance of un-transmitted packets and excessive battery units. In the numerical analysis we use $x = y = 1$, but our model can accommodate any x, y.

Computing the Optimal Control. Solving the Dynamic Programming recursion equations results in the optimal controls for (m, p^*) for all states $(b, \pi, i, mode)$. The DP terminates when the buffer empties $b = 0$ or the battery empties $\pi = 0$. Any policy that does not terminate in finite time will incur an infinite (backlog) cost. However, being ON and using any p, we can either empty the buffer with positive probability or finish our battery in finite time. Therefore, there exists a stationary optimal control solution, obtainable by value iteration. [19].

Choice of Initial Battery. Having found the optimal control for various values of the initial battery π_0, we can take an additional step and optimize the choice of initial battery required, to minimize the total cost. This is an important decision that determines the design and engineering of the battery to be put in the device: it should be large for sufficiently long operation and should also conform to practical size/form/cost constraints. To express this goal in the context of the DP formulation, we want to choose π_0^* s.t.:

$$\min_{\pi_0} \sum_{j \in I} q_{ij} J(b, \pi_0, i, mode) \tag{5}$$

IV. ADDITIONAL OPERATIONAL CONSIDERATIONS

In this section, we discuss how the general formulation and methodology can be appropriately adjusted to address various operational scenarios. First, we present two special cases of the general model; then, we mention two possible extensions.

A. Negligible Wake-up Cost

In some operational scenarios, the power cost for transitions between operating modes may be negligible compared to the other costs, i.e. $P_{on \to sleep} \simeq P_{sleep \to on} \simeq 0$. In these cases, the general model can be significantly simplified. When the radio decides not to transmit ($p = 0$), it can also switch to SLEEP, in order to avoid spending operating power P_{on}. Therefore we can eliminate the radio $mode$ (ON or SLEEP) from the system state, and the control m from the control variables. $p = 0$ now automatically means that the radio is also in $mode = SLEEP$. This was not the case in the general model, where the radio could decide to stay ON without transmitting ($p = 0$), in order to avoid the transition cost P_{tr}.

The state of the system then becomes simply (b, π, i) and the only control is the transmission power p. $p = 0$ automatically means that the radio is in SLEEP mode: there is no operating cost ($P_{on} 1_{p=0} = 0$) and there is no stress on the channel ($\Psi(p = 0, i)$ should be forced by definition to be 0). The recursive equations then become simpler:

$$\begin{aligned} J(b, \pi, i) = \min_{p=0,1,\ldots,\pi - P_{on}} \{ & B(b) + P_{on} 1_{\{p>0\}} + p \\ & + \Psi(p, i) + s(p, i) \sum_{j \in I} q_{ij} J(b - 1, \pi - p - P_{on} 1_{\{p>0\}}, j) \\ & + [1 - s(p, i)] \sum_{j \in I} q_{ij} J(b, \pi - p - P_{on} 1_{\{p>0\}}, j) \} \end{aligned} \tag{6}$$

for $b \in \{0, 1, 2, ..., b_0\}$, $\pi \in \{0, 1, 2, 3, ..., \pi_0\}$, $i \in I$.

B. Device Always Powered-On

Another special case is when the device stays ON, i.e. we disable the SLEEP mode. This scenario may arise when devices are power-on always or for a long period of interest. Even more interestingly, it is also relevant when the device has a longer duty cycle, consisting of ON and OFF periods, but we focus and optimize specifically the ON part of the duty cycle. In this case, the recursive equations get further simplified:

$$\begin{aligned} J(b, \pi, i) = \min_{p=0,1,\ldots,\pi - P_{on}} \{ & B(b) + [P_{on} + p + \Psi(p, i)] \\ & + s(p, i) \sum_{j \in I} q_{ij} J(b - 1, \pi - p - P_{on}, j) \\ & + [1 - s(p, i)] \sum_{j \in I} q_{ij} J(b, \pi - p - P_{on}, j) \} \end{aligned} \tag{7}$$

C. Packet Arrivals

So far, we consider that a certain amount of packets is already stored in the buffer and need to be transmitted. However, data may continuously arrive at the node, caused by continuous activity of the user, e.g. a VoIP call. This scenario can be easily included in the DP formulation by incorporating a live source(s) generating packets that arrive to the buffer. This will add a second source of uncertainty to the system, in addition to the interference.

D. Controlling Multiple Components of a Device

A portable/handheld device consists of several components, including the processor, the memory and the wireless radio (which has been our focus so far). Power-management of the device as a whole is also possible in the DP framework: it should consider all components, their power consumption characteristics and the interactions between them. A power analysis of the specific device of interest is important in order to understand how power is allocated to the various components and where there is room for optimization. E.g., 802.11 radios have very different power consumption characteristics than the motes used in sensor networks.

E. Network-Wide Power Management

Network of nodes. A natural next step, after the power management for a single node, is the maximization of the lifetime of an ad-hoc network consisting of multiple such nodes. Much effort has already been put to this direction from a system and protocol perspective in the context of ad-hoc and sensor networks [15], [16], [17]. As an extension of the current work, we plan to address the question of maximizing the lifetime of a battery-operated network, within the dynamic programming framework. Interesting issues include (i) studying battery savings in conjunction with topology and routing and (ii) designing practical distributed heuristics that approximate the global optimum. This class of problems bears similarities with recent work on networks-on-chips [18].

Responsive interference. A different aspect that arises when we consider many nodes is that of responsive vs. markovian interference. So far, we have modeled the channel as a Markov chain. This modeling assumption was a methodological step that allowed us to abstract and summarize the interference caused by a large number of radios into a single "background" interference. In practice, this Markov chain will have a large number of states (corresponding e.g. to nodes starting and finishing transmissions) or the channel may not even be Markovian (if many nodes implement power control). As for the first concern, our approach is clearly able to address any markovian channel, even with a large number of states, at the cost of higher computational complexity. As for the second concern, preliminary simulations in a responsive interference environment provide a sanity check.

V. SPECIAL CASE I: SENSOR ALWAYS ON, CONTROL TRANSMISSION POWER

Problem Setup. We now consider a special case, which is interesting in itself and also allows us to highlight fundamental tradeoffs of the general problem. In particular, and in order to highlight the relation between π_0 and b_0, we consider the following problem setup. First, we omit the backlog cost $B(b)$. Second, we ignore the interference i and capture the channel behavior using a probability of success $s(p)$ which is an increasing, convex function of p, e.g. $s(p) = p/(p+i)$. (In a sense, the interference i is still captured as a parameter in $s(p)$.) Third, we are always ON, paying P_{on} in every time slot, and we are never idle ($p > 0$).

The problem now is as follows. We want to send b_0 packets over a channel with success probability $s(p)$.

1) How much battery π_0^* do we need? What is the relation between π_0 and b_0? Is it linear? How does is it affected by the channel behavior $s(p)$?

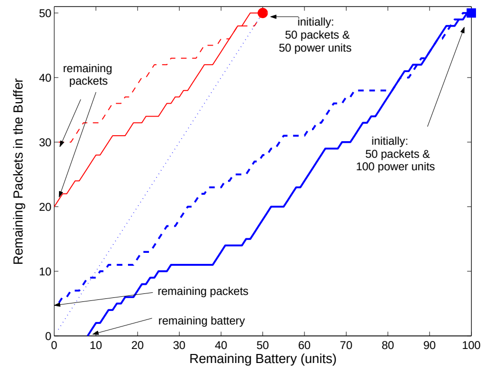

Fig. 2. Four examples of the system evolution (b, π). Starting at (b_0, π_0) (top-right side of the graph) and using some power at every time slot, we gradually empty the battery and/or the buffer (moving towards the bottom-left side of the graph). We terminate either when we empty the buffer ($b = 0$) or when we run out of battery ($p = 0$).

2) What is the optimal transmission policy p^* ? How does it depend on the remaining battery π, the current backlog b and the operational cost P_{on}/p?

The DP formulation is now simply:

$$J(b, \pi) = \min_{p=1,\ldots,\pi-P_{on}} \{W(p + P_{on}) + s(p)J(b-1, \pi - p - P_{on}) + [1 - s(p)]J(b, \pi - p - P_{on})\} \quad (8)$$

$$J(0, \pi) = W\pi, \ \ J(b, 0) = b \quad (9)$$

$$J(b) = \min_{\pi_0} J(b, \pi_0) \quad (10)$$

The weight W in front of the power-related costs indicates how much we value power vs. packet delivery and it strongly affects the optimal policy.

Intuition. The example illustrated in Fig. 2 gives some intuition. Let's say we want to send $b_0 = 50$ packets. Assume the optimal power policy is to use 1 unit for transmission every time and ignore P_{on}. If the channel were perfect $s = 1$, we would need exactly one unit of power per packet: $\pi_0 = b_0$. In Fig. 2, the system would evolve from $(b_0 = 50, \pi_0 = 50)$, across the straight dotted line toward (0,0), transmitting all packets and wasting no battery.

If $s < 1$ when using 1 unit power to transmit, some packets are lost and additional power is needed to retransmit them: $\pi_0 > b_0$. E.g. for a Bernoulli channel w.p. $s < 1$, we need on average $\pi_0 = b_0/s \geq b_0$ units of power. For example, in Fig. 2, the top thin lines show the sample paths exercising optimal control starting at $(b_0 = 50, \pi_0 = 50)$ and assuming $s(p = 1) = 0.1$; the system always terminates on the y-axis $(b > 0, \pi = 0)$, even if we give a lot of importance (large W) to power. However, if we starting at higher initial power $(b_0 = 50, \pi_0 = 100)$, we are able to terminate at the y- or at the x-axis, by appropriately tuning W.

For a general success probability $s(p)$, the optimal choice of initial battery π_0^* depends on the shape of $s(p)$. The DP

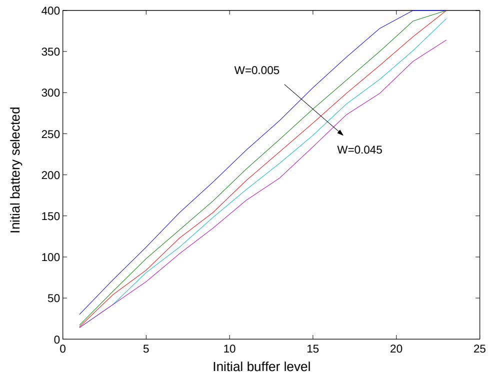

Fig. 3. Optimal initial battery π_0^* increases with the number of packets to be transmitted b_0, depending on the weight W. (Keeping i, and thus $s(p)$, the same.)

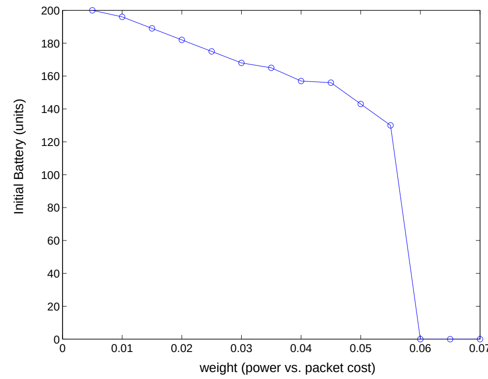

Fig. 5. Optimal initial battery π_0^*, depending on the weight W (relative importance of power vs. remaining packets).

recursive equations compute the optimal control (π_0^*, p^*). After choosing the right π_0^*, we let the system evolve from (b_0, π_0^*) until it hits $b = 0$ or $\pi = 0$, using the optimal control p^*. Depending on how much we value packets vs. power (captured by our choice of weight W), we can affect the optimal policy p^* and make the system "hit" one of the two axes. Ideally, we would like to hit point (0,0), which means that we used just enough power to transmit the packets.

Structural Properties of the Optimal Policy. We now numerically compute the optimal policy (π_0^*, p^*) and comment on its observed structural properties. The power unit in all figures is the step in Tx power $p = 1, 2, ...P_{max}$ (i.e. we normalize w.r.t. $p_{min} \doteq 1$)

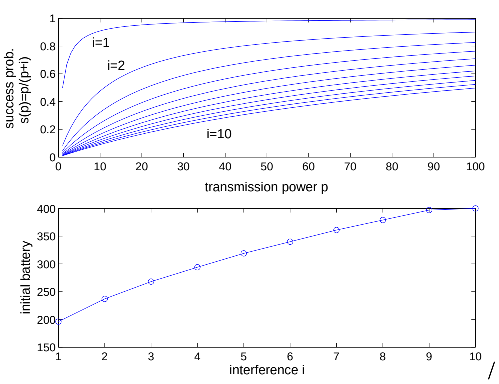

Fig. 4. *Effect of Channel Volatility.* Increasing i (a) the success probability $s(p) = p/(p + i)$ becomes more volatile and (b) larger initial battery π_0 is required.

Fig. 3 confirms the intuition that more initial power π_0^* is needed to send more packets b_0; the slope depends on the weight W, the relative importance of battery vs. remaining packets in the buffer. Furthermore, π_0^* depends on the volatility of the channel, as shown in Fig. 4: as i increases, $s(p)$ becomes more volatile and we need larger π_0 to accommodate the channel fluctuations even for the same b_0 and W.

Fig. 5 shows in more detail the effect of the weight W. The larger the weight on power, the more conservative the optimal transmission policy p^* , the less initial π_0 we need; this is captured by the linear decreasing part in Fig. 5. At the extreme, where we choose W such that $W(p + P_{on}) \geq s(p), \forall p$, we indicate that we value power spent in one time slot more than the delivery of one packet $s(p)$. In that case it doesn't even worth it to try to send anything, which is captured, in Fig. 5, by the threshold effect and the sharp decrease down to $\pi_0 = 0$.

In summary, regarding the optimal initial battery π_0^*:

- $\pi_0^* \uparrow$ as $b_0 \uparrow$
- $\pi_0^* \uparrow$ as $i \uparrow$
- $\pi_0^* \downarrow$ as $W \uparrow$. $\pi_0^* = 0$ if $W(p + P_{on}) > s(p)\ \forall p$

Let us now discuss the properties of the optimal transmission policy $p*$. Fig. 6 shows p^* as a function of the remaining battery π, for one packet ($b = 1$) and for various values of P_{on}. In transmitting one packet we face the following dilemma. On one hand, we can use low Tx power p_l, thus have low prob. of success $s(p_l)$, and spend several (say k) slots and $k(p_l + P_{on})$ units of power before we succeed. On the other hand, we can spend a large p_h and get the packet through with high prob. in one slot, paying $p_h + P_{on}$. If P_{on} is small, then we can afford paying it multiple times. As $P_{on} \uparrow$ (moving from the top to the bottom plots of Fig.6), we cannot afford paying P_{on} multiple times and we become more aggressive: we use high Tx power in order to finish as soon as possible. The chosen number of time slots k is constrained by whether there is enough battery for all k of them. This explains the triangular shape along the π-axis: as more power is available, there are more transmission opportunities which allow us to decrease p^* and spread it over more slots. The sharp decreases in Fig.6, happens when $\pi \uparrow$ so that we get one more opportunity to transmit (say from k to $k + 1$ times); thus the periodic triangular shape. When more battery is available, we have ample transmission opportunities to transmit and we become less stressed: the triangular shape is repeated, but the fluctuations have lower peaks, as we move to right of the π-axis.

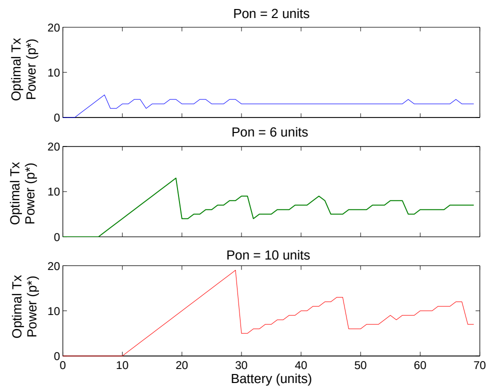

Fig. 6. *Effect of P_{on}*. Optimal transmission power p^*, for transmitting one packet ($b = 1$), as a function of the battery reserves π and the operating cost P_{on} to be on.

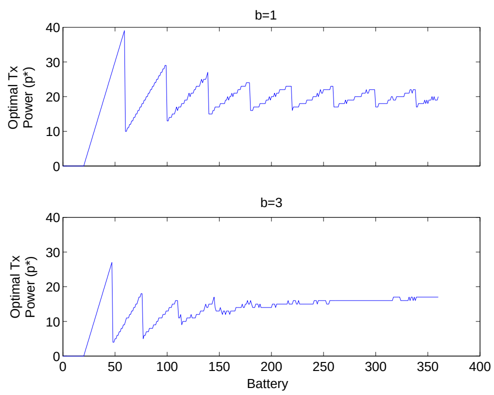

Fig. 7. *Effect of battery π on the optimal policy p^**. Optimal transmission power p^* for transmitting b packets, as a function of battery reserves π.

The exact same tradeoff also explains the triangular shape in Fig.7. As more battery π is available, the choice of p^*, (and thus $s(p)$ and eventually number of slots) is less critical. Thus the triangular shape is more pronounced at the left side, where we run short of battery. In the bottom plot of Fig.7, we show the case for multiple packets: the oscillations are less pronounced and they have smaller period. This is because, we cannot be aggressive for the first packets as we have to save some battery for later packets as well. The fluctuations again disappear as we move to the right (large amount of battery).

Understanding the structural properties of the optimal policy is important in order to design practical heuristics. Preliminary results have shown that heuristics mimicking these structural properties, achieve near-optimal performance at lower implementation complexity.

VI. SPECIAL CASE II: CONTROLLING BOTH TRANSMISSION POWER AND SLEEP MODE

We now extend the previous section by controlling not only the transmission power but also the operation mode of the radio. For example, this might be useful when the channel interference is high; in order to avoid spending a large amount of operating power, the radio may choose to go into SLEEP mode. While waiting for the interference to decrease, the radio incurs a backlog cost due to the packets remaining in the buffer; this pressure will eventually force the radio to transmit. The tradeoff is now between power savings and backlog cost.

In order to better highlight this core tradeoff, we make the following choices within the general model of section III. First, we consider the i.i.d. case for the channel model: a low interference state (i_l) w.p. p_l, and a high interference (i_h) state w.p. p_h. The probability of successful transmission remains $s(p,i) = p/(p+i)$. For the numerical examples in this section, we use $p_l = p_h = 1/2$, unless stated otherwise. The state now includes the backlog b, the battery π, and the interference level i. We now include in our control the option of $p = 0$ (SLEEP mode), during which no operating cost (P_{on}) is incurred. For a fair comparison with section V, we model the power and backlog costs similarly. As in the previous section, the weight W in front of the power-related costs indicates how much we value power vs. packet delivery and it strongly affects the optimal policy. The recursive equations become as follows (notice the dependence on the time slot n):

$$\begin{aligned} J_n(b,\pi,i) = & \min_{p=0,\dots,\pi-P_{on}} \{W(p + P_{on}1_{p>0}) \\ & + s(p,i) \sum_{level=l,h} p_{level} J_{n+1}(b-1, \pi - p - P_{on}1_{p>0}, i_{level}) \\ & + [1 - s(p)] \sum_{level=l,h} p_{level} J_{n+1}(b, \pi - p - P_{on}1_{p>0}, i_{level})\} \end{aligned} \tag{11}$$

$$J_N(0,\pi,i) = W\pi, \;\; J_N(b,0,i) = b \tag{12}$$

$$J(b) = \min_{\pi_0} J(b,\pi_0) \tag{13}$$

Fig. 8 shows the effect of P_{on} on p^* to transmit one packet with 17 time slots left before termination. We observe that when the channel interference is high, the optimal policy chooses to not transmit and enters in SLEEP mode. When the channel interference is low, the shape of the optimal policy shows similar dependencies on P_{on} as in Fig. 6. When P_{on} is low, we can afford to transmit frequently and transmit at lower power. Conversely, when P_{on} is high, we prefer to transmit at high power only a few times. Unlike the previous section, when the channel is good, transmission power increases with the battery reserves. This is due to the dependence on time in the DP formulation; the radio spends more power to ensure successful transmission and avoid backlog costs.

This backlog pressure is better exemplified in Fig. 9. With termination in 2 time slots, we must aggressively transmit to avoid the cost of unsent packets. Unlike Fig.8, the radio transmits even when interference is high.

These trends are also apparent in Fig. 10. Again, the dependence on time in the DP formulation affects the

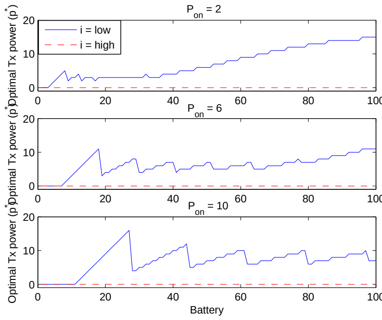

Fig. 8. *Effect of P_{on}.* Optimal transmission power p^* as a function of the battery reserves π and P_{on} within 17 time slots.

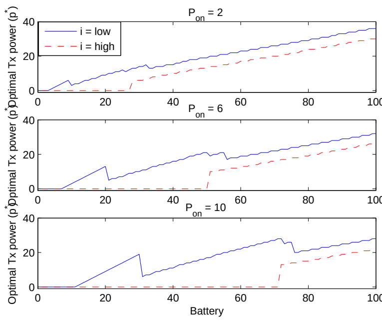

Fig. 9. *Effect of P_{on}.* Optimal transmission power p^* as a function of the battery reserves π and P_{on} with 2 time slots.

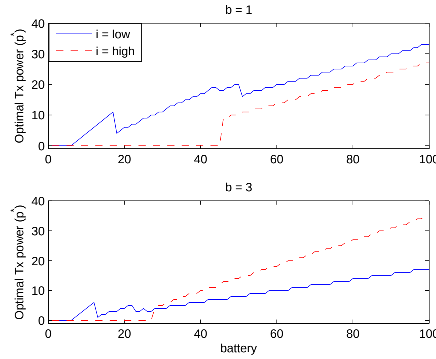

Fig. 10. *Effect of battery.* Optimal transmission power p^* as a function of battery reserves, for transmitting b packets within 2 time slots.

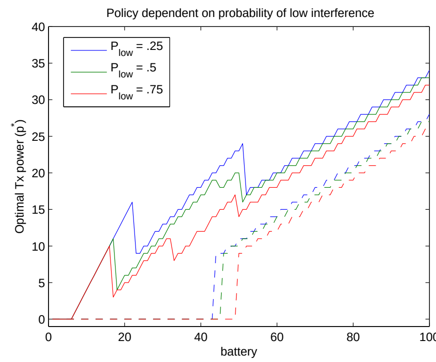

Fig. 11. *Effect of p_l and p_h on the optimal policy.* Optimal transmission power p^* for transmitting 1 packet with 2 time slots until termination, when the probability of each channel state is varied.

optimal policy. When there is only 1 packet to transmit, the radio aggressively transmits even when interference is high. When there are 3 packets to transmit, the radio becomes even more aggressive, because there are only 2 time slots left for transmitting all 3 packets. When there are many time slots, the optimal policy is more relaxed: it avoids transmitting when the channel is bad, even with a large number of packets in the buffer.

Fig. 11 shows the effect of the channel state probabilities, p_l and p_h. When the probability of the state with low interference is high (e.g. $p_l = .75$), the radio can use a conservative transmission policy. When this probability is lower (e.g. $p_l = .25$), the radio must be more aggressive and take advantage of the low interference state to ensure successful packet transmission before re-entering the more common high interference state. Fig. 11 shows these effects at two times slots before termination. When there are more time slots to go, the same observations hold, but the backlog pressure is lower; therefore, the radio can afford to enter in SLEEP mode when the channel is bad.

VII. Comparison of Transmission Policies

In this section, we compare the two optimal transmission policies from sections V and VI, with each other as well as with two benchmarks policies (constant power and constant signal-to-interference ratio). More specifically, the policies under comparison are the following:

1) The optimal policy of section V. This policy is unaware of the channel interference level and is only aware of the average interference.
2) The optimal policy of section VI. This is a channel aware policy: we assume that the radio can probe the channel and has knowledge of the interference level.
3) The commonly used ”constant power” benchmark. It is channel unaware and uses the same power to transmit each packet, so $p^* = \pi/b$.
4) Another commonly used benchmark is the policy that achieves constant signal-to-interference ratio (or constant SIR). Because the constant SIR benchmark is channel aware, in order to make a fair comparison, we consider a constant power benchmark that uses power equal to the average transmission power used by constant SIR.

We performed numerical simulations with the following setup. The channel was assumed i.i.d. with $p_l = p_h = 1/2$. The initial buffer level was $b = 10$. We varied the initial battery reserve, π, and measured the average number of successful packet transmissions achieved by each policy. For the same initial battery reserve, the larger the number of successful transmitted packets, the better the performance of the policy. Fig. 12 and Fig. 13 show the results of this comparison for two sets of interference levels.

In Fig. 12, $i_h/i_l = 50$. The channel unaware policy outperforms the constant power benchmark: it successfully transmitting over twice the number of packets. Also, our channel aware policy successfully transmits twice as many packets as the constant SIR policy. The channel unaware policy barely outperforms the SIR benchmark, because the

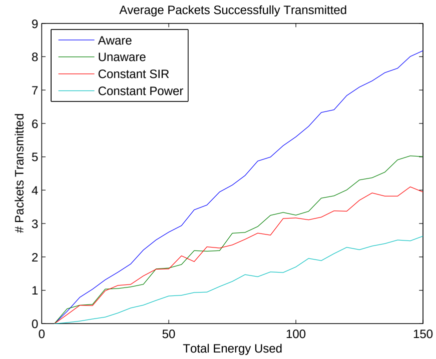

Fig. 12. *Comparison of different transmission policies.* Beginning with 25 packets in the buffer, the total number of packets successfully transmitted as a function of the total energy used. $i_l = 1$ and $i_h = 50$.

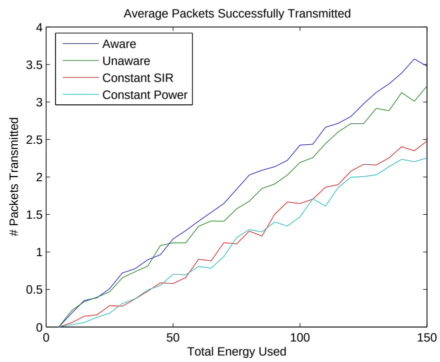

Fig. 13. *Comparison of different transmission policies.* Beginning with 25 packets in the buffer, the total number of packets successfully transmitted as a function of the total energy used. $i_l = 10$ and $i_h = 50$.

latter has the advantage of using information about the channel interference.

In Fig. 12, $i_h/i_l = 5$, i.e. the difference between the high and low interference levels is smaller. Because the interference levels are similar, the benefit from knowing the accurate interference level, as opposed to the average interference level is less valuable. Therefore, the benefit from using channel aware, instead of channel unaware, policies is smaller. Still, the channel aware policies outperform the channel unaware policies. We can also see that the transmission policies we have presented in this paper continue to outperform the two benchmark policies.

VIII. CONCLUSION

In this paper, we studied power management at the radio of a battery-operated portable device, so as to transmit a certain amount of information while preserving the battery. We also showed how to optimally choose the initial battery given the various design costs. The main contribution of the paper is the formulation of this class of problems in the dynamic programming framework. Furthermore, we considered both aspects of (i) dynamic power management and (ii) transmission power control, and we explored the relative energy savings from these two, for a range of conditions (captured by $p : P_{on}$). As an illustration of our approach, we analyzed two insightful special cases, that allowed us to highlight fundamental tradeoffs and characteristics of the optimal policy. We are currently working on more extensive simulations and heuristics design. We hope that this approach will be useful for characterizing the performance limits of battery-constrained portable devices and for designing and evaluating practical heuristics.

REFERENCES

[1] L.Benini, G.De Micheli, “Dynamic Power Management: design techniques and CAd tools,” Kluwer, 1997.

[2] L. Benini, G. Paleologo, A. Bogliolo, and G. D. Micheli, “Policy optimization for dynamic power management,” IEEE Trans. Computer-Aided Design, vol. 18, pp. 813-833, June 1999.

[3] A.Karlin, M.Manesse, L.McGeoch and S. Owicki, “Competitive Randomized Algorithms for Nonuniform problems”, *Algorithmica*, pp.542-571, 1994.

[4] M.Srivastava, a.Chanrakasan, R. Brodersen, “Predictive system shutdown and oter architectural techniques for energy efficient programmable computation”, *in IEEE Trans. on VLSI Systems*, Vol.4, no.1, pp.42-55, March. 1996

[5] N.Bambos and S.Kandukuri, ”Power Controlled Multiple Access (PCMA) in Wireless Communication Networks,” *INFOCOM 2002*, pp. 368-395.

[6] E.Shih, P.Bahl, M.Sinclair,“Wake on Wireless: An event Driven Energy Saving Strategy for Battery Operated Devices”, *in Proc. of ACM Mobicom 2002*, Sept. 2002, Atlanta.

[7] S.Agarwal, S.Krishnamurthy, r. Katz, S.Dao, “Distributed Power Control in Ad-Hoc Wireless Networks”, *in Proc. of IEEE PIMRC’01*, pp59-66, 2001.

[8] T.Simunic, L.Benini, P.Glynn, G. De Micheli, “Dynamic Power Management for Portable Systems”, *in Proc. of ACM Mobicom 2000*.

[9] D.Qiao, S.Choi, A.Jain, K.Shin, “MiSer: An optimal Low-Energy Transmission Strategy for IEEE 802.11a/h”, *in Proc. of ACM Mobicom 2003*, Sept. 2003, San Diego, CA, USA.

[10] D.Qiao, S.Choi, A.Soomro, K.Shin, “Energy-Efficient PCF Operation of IEEE 802.11a LAN”, *in Proc. of ACM Mobicom 2002*, June 2002, New York, USA.

[11] *IEEE 802.11. Part II: Wireless LAN Medium Access control (MAC) and Physical Layer (PHY) Specifications.* IEEE. Aug.1999.

[12] *IEEE 802.11h/D2.2. PART II: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Spectrum and Transmit Power Extensions in the 5GHz band in Europe.* Supplement to 802.11 IEEE Standard-1999, D 2.2, Sept. 2002.

[13] V. Raghunathan, C. Schurgers, S. Park, and M. Srivastava, “Energy-aware wireless sensor networks”, *IEEE Signal Processing Magazine* 19, 2, 40-50, March 2002.

[14] C. Hsin , M. Liu, “Network coverage using low duty-cycled sensors: random & coordinated sleep algorithms”, in Proc. 3rd IPSN 2004.

[15] B.Chen, K. Jamieson, H.Balakrishnan, R. Morris, “Span: Energy Efficient Coordination for Topology Maintenance in Ad Hoc Networks”, *Mobicom 2001*.

[16] Y.Xu, J.Heidemann, D. Estrin, “Geography-informed Energy Conservation for Ad-hoc Routing”, *in Proc. of ACM MobiCom 2001*.

[17] Ramanathan and R. Rosales-Hain, ”Topology control of multihop wireless networks using transmit power adjustment”, *in Proc. of IEEE Infocom 2000*.

[18] T.Simunic, S.Boyd, P.Glynn, “Managing power consumption in networks on chips”, *in IEEE Trans. on VLSI*, Vol.121(1), pp. 96-107, Jan.2004.

[19] D.Bertsekas, ”Dynamic Programming and Optimal Control”, Athena Scientific, 2nd Edition.

[20] Piyush Gupta and P. R. Kumar, “Towards an Information Theory of Large Networks: An Achievable Rate Region,” *IEEE Trans. on Information Theory*, vol. 49, no. 8, pp. 1877–1894, August 2003.