Energy Efficient DSL via Heterogeneous Sleeping States: Optimization Structures and Operation Guidelines

Ioannis Kamitsos[1], Paschalis Tsiaflakis[2], Ken Kerpez[3], Sangtae Ha[1], and Mung Chiang[1]
[1]Princeton University, USA, [2]ESAT-SCD, KU Leuven, Belgium, [3]ASSIA, Inc., USA

***Abstract*—Switching off a DSL line to a low-power sleeping state is becoming an important method to enhance energy efficiency of DSL broadband access networks. Although a low-power (L2) state and an off (L3) state are already defined in DSL standards, they have not been fully exploited due to concern about the resulting time-varying crosstalk. In this paper, we develop a method to optimally choose the appropriate sleeping state based on the modem's switching cost characteristics and the power consumption incurred during each sleeping state. We further develop the optimal policy for switching to the appropriate sleeping state, allowing transitions between operating modes with different power and transmit rate. Numerical results on our realistic DSL simulator show that more flexibility in control, introduced by our approach, improves the energy-delay Pareto optimal tradeoff, and results in a more energy efficient and stable DSL system compared to the existing power saving policies.**

I. Introduction

Given the global penetration of Digital Subscriber Line (DSL) technology, significant research has been devoted to improving its energy efficiency [1][2]. In addition, the EU Code of Conduct (CoC) mandates energy efficiency for DSL equipment and encourages the use of low-power states. It is becoming necessary for equipment manufacturers to comply with the CoC. However, transitions to and from these low-power "sleeping" states must consider both service delay and stability. Specifically, putting a DSL line in a low-power sleeping state saves energy but at the same time reduces its service rate and thus increases service delay. In addition, switching to and from low-power sleeping states comes with a switching cost in energy and time. Even more importantly, and a critical issue in practice, switching generates time-varying crosstalk to other DSL lines [3], because of the presence of electromagnetic coupling among different lines within the same cable bundle. This network instability should be avoided as much as possible.

This paper derives the optimal method of controlling low-power states for DSL lines, while improving the overall stability of the DSL network, and analyzes the resulting performance benefits and tradeoffs. A general discussion of the problem space and initial results were recently presented in [4]. In this paper we substantially expand the previous theory and results:

1) We optimize the use of **multiple** heterogeneous power states: L0 (full power), L2 (low power), and L3 (off). We provide a method to optimally decide which sleeping state (L2 vs L3) to use. L2 has a non-zero power consumption but a lower switching cost, whereas L3 has zero power consumption but a higher switching cost. Taking into account this tradeoff, we calculate the optimal thresholds for changing the sleeping state from L3 to L2 and vice versa. Fig. 1 shows the power consumption and switching cost incurred during transitions between the three different power states in DSL modems.

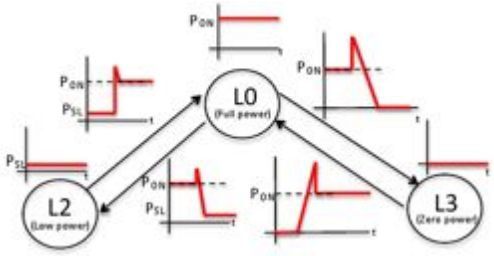


Fig. 1. Transitions between different power states (L0, L2, and L3) in DSL modems. The power consumption at L2 is non zero but the switching cost for L2/ON transitions is lower both in energy and time. The power consumption at L3 is zero but the switching cost for L3/ON transitions is larger both in energy and time. The power consumed during L0 and L2 is denoted as P_{ON} and P_{SL} respectively.

2) We propose an analytic framework for an optimal sleeping policy that **gradually** transitions a line between operating modes of different power consumption and transmit rate. We prove the structural properties of the optimal policy, which allows the modem to switch between different operating modes, and does not require a full transition to the sleeping state (L2 or L3) before switching back on.
3) We quantify the benefits that arise from (a) the increased number of operating modes and (b) improving the flexibility in switching between the different operating modes. We demonstrate that these two degrees of freedom result in significant improvements in terms of energy-delay Pareto optimal tradeoff, and noise impact on uncoordinated DSL lines.
4) The Broadband Forum (BBF), which is the central body for next-generation broadband specifications, has developed over 200 technical specifications for broadband, many of which are in widespread use. The BBF now has an active project on DSL Energy Efficiency and specifies a procedure to use the L2 mode in DSL modems. It recommends the DSL line to switch off to L2 via 9 intermediate states with power consumption 1dB apart, spending 127s in each intermediate state. While staying

TABLE I
MAIN NOTATION

Symbol	Meaning [Units]
P_{ON}	Power consumed when DSL modem is ON [Watts]
B	Buffer size [Number of jobs]
$i=(W^1,W^2,Q)$	i-th state in the state space
W^1	Defines whether the DSL modem is in L2 or not
W^2	Defines whether the DSL modem is in L3 or not
Q	Queue length [Number of jobs]
R	Service rate [jobs/s]
λ	Arrival rate [jobs/s]
$\mathbf{p}(i)$	policy for state i
P_{SL}	power consumption at L2
τ_{off}	Time spent in each of the intermediate states defined by BBF guidelines
E^1_{ch}	Energy consumed on transitions ON$\rightarrow$ L2 and L2$\rightarrow$ ON [kWh]
E^2_{ch}	Energy consumed on transitions ON$\rightarrow$ L3 and L3$\rightarrow$ ON [kWh]
τ_1	Time required for DSL modem to wake up from L2 [s]
τ_2	Time required for DSL modem to wake up from L3 [s]
$M^{\mathbf{p}(i)}_{i\rightarrow j}$	Transition probability of moving from state i to state j under policy $\mathbf{p}(i)$
$V(i)$	Average cost of state i
$g(i,p_1(i),p_2(i))$	Cost per transition of state i under policy $\{p_1(i),p_2(i)\}$
δ	Switching time cost coefficient [J/job]
r	Congestion cost coefficient [W/job]
α	Discount factor in dynamic programming

at each state for this long duration assists stability, we show that the energy savings are very poor. In particular, we demonstrate that our sleeping policy results in (a) higher energy savings even when the BBF's procedure is modified to stay in each intermediate state a very short amount of time for better energy savings (e.g., 1.27s, 100 times smaller than the BBF recommendation), and (b) a more stable operation even when our approach turns the modem off more aggressively via only 6 operating modes.

Other recent work on the exploitation of DSL low-power states for obtaining energy savings is presented in [9], which focuses on selecting the rate at the low-power states by analyzing traffic-dependent energy savings.

II. Optimizing over L2 vs L3 sleeping states

A. Model and Analysis

We assume that each modem has a finite buffer B and three heterogeneous modes of operation: ON mode and two sleeping modes (L2 and L3). The power that the modem consumes when ON is P_{ON}, whereas the power consumption at L2 is P_{SL}, with $0<P_{SL}<P_{ON}$. The power consumption at L3 is considered to be 0. The transmit rate in ON mode is R, whereas the transmit rate at L2 and L3 is $R_{SL}>0$ and 0 respectively. To simplify the analytic derivations we assume that $R_{SL}\approx 0$ without loss of generality, given that $R_{SL}<<R$. A summary of the notation can be found in Table I.

We also assume that jobs arrive to the modem according to a Poisson process with rate λ. The job sizes are exponentially distributed with mean $1/R$. Based on the memoryless property of the arrival process we formulate our problem as a Markov Decision Process (MDP) problem.

The energy consumed in ON/L2 and ON/L3 switchings is E^1_{ch} and E^2_{ch} respectively, with $E^2_{ch}>E^1_{ch}$. Also waking up from L2 and L3 takes time τ_1 and τ_2 respectively with $\tau_2>\tau_1$. We assume that there is no transition between L2 and L3. This gives rise to the following tradeoff: *Whenever we want to turn off the modem, do we switch it to L2 with non-zero power consumption but with low switching cost, or do we switch it to L3 mode with zero power consumption but with higher switching cost?*

Each state in the state space Ω is characterized by the triplet (W^1,W^2,Q), where $W^1,W^2\in\{0,1\}$ denotes whether the modem was in L2 or L3 respectively during the previous time interval, and Q is the number of jobs in the queue. If $W^1=W^2=0$ this means that the modem was ON during the previous time interval.

Let $\mathbf{p}(i)=\{p_1(i),p_2(i)\}$ be the optimal policy of state i, where the component $p_1(i)\in\{0,1\}$ defines whether it is optimal or not to put the modem to L2 in the next state, and the same rule applies to $p_2(i)\in\{0,1\}$ for L3. If $p_1(i)=p_2(i)=0$, then it is optimal to have the modem ON at the next state. A case with $p_1(i)=p_2(i)=1$ is not valid.

The states evolve as follows:

$$W^1(t+\tau)=p_1\left(W^1(t),W^2(t),Q(t)\right), \tag{1}$$

$$W^2(t+\tau)=p_2\left(W^1(t),W^2(t),Q(t)\right), \tag{2}$$

$$Q(t+\tau)=\begin{cases}Q(t)+1, & \text{if an arrival occurs,}\\ Q(t)-1, & \text{if a departure occurs.}\end{cases} \tag{3}$$

Note that t denotes the current time and $t+\tau$ refers to the time of the next event.

The transition probabilities from state (W^1,W^2,Q) are

$$Pr[(W^1,W^2,Q)\rightarrow(p_1(W^1,W^2,Q),p_2(W^1,W^2,Q),Q+1)] = \frac{\lambda}{\lambda+R(1-p_1(W^1,W^2,Q)-p_2(W^1,W^2,Q))} \tag{4}$$

$$Pr[(W^1,W^2,Q)\rightarrow(p_1(W^1,W^2,Q),p_2(W^1,W^2,Q),Q-1)] = \frac{R(1-p_1(W^1,W^2,Q)-p_2(W^1,W^2,Q))}{\lambda+R(1-p_1(W^1,W^2,Q)-p_2(W^1,W^2,Q))} \tag{5}$$

The cost per transition consists of the running cost $g_1(i,p_1(i),p_2(i))$:

$$g_1(i,p_1(i),p_2(i))=(1-p_1(i)-p_2(i))P_{ON}+p_1(i)P_{SL}+rQ_i \tag{6}$$

and the switching cost $g_2(i,p_1(i),p_2(i))$:

$$g_2(i,p_1(i),p_2(i))=|p_1(i)-W^1_i|E^1_{ch}+|p_2(i)-W^2_i|E^2_{ch}+ +\delta|p_1(i)-W^1_i|W^1_i\lambda\tau_1+\delta|p_2(i)-W^2_i|W^2_i\lambda\tau_2. \tag{7}$$

Note that we also take into account the time it takes for the modem to wake up from L3 or L2, which we express in terms of the $\lambda\tau_1$ or $\lambda\tau_2$ jobs that accumulate in the buffer during the modem wake up process.

To analyze the energy-performance tradeoff, we aim at minimizing the average discounted sum of costs

$$V(i) = \min_{p_1(i),p_2(i)} \{A(p_1(i),p_2(i)) + B(p_1(i),p_2(i))\}, \quad (8)$$

where

$$A(p_1(i),p_2(i)) = \frac{g_1(i,p_1(i),p_2(i))}{\beta+v} + g_2(i,p_1(i),p_2(i)), \quad (9)$$

and

$$B(p_1(i),p_2(i)) = \alpha \sum_{j\in\Omega} \hat{M}_{i\to j}^{(p_1(i),p_2(i))} V(j). \quad (10)$$

where v is the uniform transition rate given by $v = \lambda + R$, and $\hat{M}_{i\to j}$ is the uniform transition probability going from state i to state j under policy $\{p_1(i),p_2(i)\}$[1]. In order to make the optimal cost finite we approximate the objective in (8) by discounting future costs with a discount factor α close to 1. This also provides robustness against uncertainty of future events. We write $\alpha = v/(\beta+v)$ and select a small $\beta << v$ in order for $\alpha \approx 1$. The optimal policy can be calculated by solving (8) using the value iteration algorithm [5].

The following two definitions are useful for the main analytic results.

Definition 1: A policy q is called *hysteretic if for all* $a \in \mathcal{F}$, $q(a,Q) = \gamma$ *implies* $q(\gamma,Q) = \gamma$.

Definition 2: [6] A hysteretic policy p is called *monotone if there are* l_α, u_α, $\alpha = 1,2,\ldots,|A|+1$, *where* $|A|$ *is the set of operating modes, with* $l_\alpha \le u_\alpha$ *for* $\alpha = 1,2,\ldots,|A|+1$, $l_\alpha \le u_{\alpha+1}$ *for* $\alpha = 1,2,\ldots|A|$, $l_\alpha = u_\alpha = 0$ *and* $l_{|A|+1} = u_{|A|+1} = \infty$, *such that for all* $(i,d) \in \Omega$, *we have*

$$q(i,d) = d, \text{if } l_d \le i < u_{d+1}$$
$$q(i,d) = q(i,d+1), \text{if } i \ge u_{d+1}$$
$$q(i,d) = q(i,d-1), \text{if } i < l_d$$

The practical interpretation of the two definitions above is as follows: when there are two levels of operation, e.g., ON and sleeping (L2 or L3), such a policy is characterized by two thresholds in terms of number of jobs in the queue - one for turning on and the other for switching off to sleeping mode. This simplifies implementation and makes the policy adaptive to the current workload.

The proof of the following theorem is provided in Appendix.

Theorem 1: **The optimal policy for the objective in (8) is hysteretic.**

B. Simulation Results

For the following simulations we consider a DSL modem with transmit power equal to 20.4dBm and data rate equal to 15.9 Mbps. These numbers are generated by using a realistic DSL simulator. We assume a rather delay insensitive application with $r = 10$J/job. The buffer size is 100 and the time required to switch the modem on or off is 40s. Similarly to [4], it is optimal to switch off the modem when the buffer drains. This section examines the previously unanswered question: under what conditions is it optimal to switch off to L2 versus to L3 and vice versa? The main parameters that affect this decision are the ratio between P_{SL} and P_{ON} ($\frac{P_{SL}}{P_{ON}}$) and the ratio between E_{ch}^2 and E_{ch}^1 ($\frac{E_{ch}^2}{E_{ch}^1}$).

In particular, Fig. 2 shows the optimal value of the ratio $\frac{P_{SL}}{P_{ON}}$, denoted as $(\frac{P_{SL}}{P_{ON}})^*$. The optimal threshold $(\frac{P_{SL}}{P_{ON}})^*$ is plotted against the traffic load for two different values of the ratio $\frac{E_{ch}^2}{E_{ch}^1}$. Note that if the actual value $\frac{P_{SL}}{P_{ON}}$ is larger than $(\frac{P_{SL}}{P_{ON}})^*$ the optimal sleeping state is L3, otherwise it is L2. For a given ratio of switching costs $\frac{E_{ch}^2}{E_{ch}^1}$, the increase in traffic

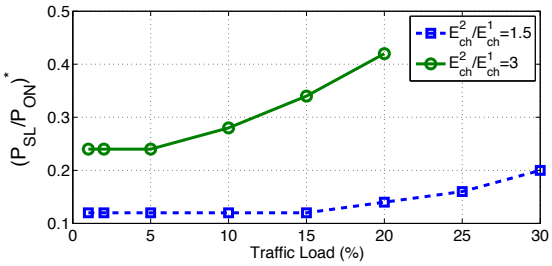


Fig. 2. Optimal ratio $(\frac{P_{SL}}{P_{ON}})^*$ for increasing traffic load. For a given ratio $\frac{E_{ch}^2}{E_{ch}^1}$, increasing the traffic load results in an increased $(\frac{P_{SL}}{P_{ON}})^*$. Given the increased probability of waking up, the modem tolerates higher values of P_{SL} in order to avoid too frequent switchings from L3. Increasing the ratio $\frac{E_{ch}^2}{E_{ch}^1}$ for a given traffic load results in the modem tolerating higher values of P_{SL} without switching off to L3.

load results in a higher $(\frac{P_{SL}}{P_{ON}})^*$. As the traffic load increases, the probability of frequent switchings between ON and the sleeping state (L2 or L3) becomes higher. Thus, switching to L2 is preferable even when the cost P_{SL} is high since it is still less expensive than frequent switchings between L3 and ON. Similarly, for a given traffic load, the modem with a larger switching costs ratio $\frac{E_{ch}^2}{E_{ch}^1}$ demonstrates a higher $(\frac{P_{SL}}{P_{ON}})^*$ to avoid the excessive energy consumption resulting from waking up from L3. For example, when the traffic load is beyond 30% for $\frac{E_{ch}^2}{E_{ch}^1} = 1.5$, or beyond 20% for $\frac{E_{ch}^2}{E_{ch}^1} = 3$, the optimal policy never switches the modem off to L3. It instead either keeps the modem always ON or switches it off to L2 when the buffer drains.

In addition, Fig. 3 shows the optimal value of the ratio $(\frac{E_{ch}^2}{E_{ch}^1})^*$ for different values of the traffic load and $\frac{P_{SL}}{P_{ON}}$. Note that if the actual value $\frac{E_{ch}^2}{E_{ch}^1}$ is larger than $(\frac{E_{ch}^2}{E_{ch}^1})^*$, the modem switches off to L2, otherwise, the modem switches off to L3. For a given value of P_{SL}, the ratio $(\frac{E_{ch}^2}{E_{ch}^1})^*$ decreases with traffic load. As the traffic load increases, the modem prefers switching off to L2 since it has high chance to turn back on. Therefore, it can tolerate consuming P_{SL} even for lower values of $\frac{E_{ch}^2}{E_{ch}^1}$. For a given traffic load, as P_{SL} increases, the optimal value of $(\frac{E_{ch}^2}{E_{ch}^1})^*$ increases. Under higher P_{SL} the modem tends

[1]Uniformization is a standard technique to write dynamic programming problems [5]

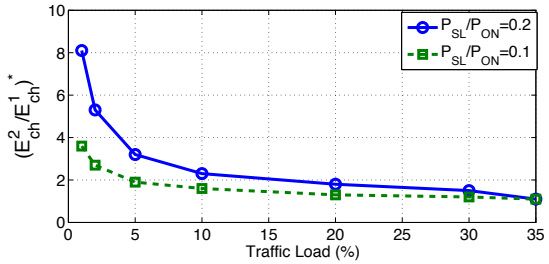


Fig. 3. Optimal ratio $(\frac{E^2_{ch}}{E^1_{ch}})^*$ for increasing traffic load. For a given ratio $\frac{P_{SL}}{P_{ON}}$, the ratio $(\frac{E^2_{ch}}{E^1_{ch}})^*$ decreases with traffic load. Given the increased probability of waking up, the modem prefers switching off to L2, even when the switching cost to L3 is lower. For a given traffic load, higher $\frac{P_{SL}}{P_{ON}}$ results in increased $(\frac{E^2_{ch}}{E^1_{ch}})^*$ since the switching cost to L3 has to be really high to prevent the modem from switching to L3.

to switch off to L3 in order to avoid consuming high P_{SL}, hence the range of values of $\frac{E^2_{ch}}{E^1_{ch}}$ where the modem switches off to L2 decreases. For traffic load beyond 35% the optimal policy never switches the modem off to L3. Instead, it either keeps it always ON or switches it off to L2 when the buffer drains.

III. Multiple Operating Modes and Gradual Transitions

A. BBF Guidelines and Our New Design

BBF guidelines [8] define a power saving procedure, where they recommend the use of L2 since waking up from L3 would create excessive crosstalk and could render other DSL lines unstable. They also suggest a gradual powering down of the DSL modem via 9 intermediate states with 1dB apart per each state and 127s spent in each state. However, turning on the modem can be instantaneous and takes place as soon as there is data to transmit.

In this section, we develop an analytic framework that takes into account the gradual turning off of the modem, but at the same time considers a gradual turning on via similar intermediate states or operating modes. We seek to find the optimal policy that controls such transitions, taking into account the fact that once a modem is in the process of waking up, it does not need to transition through all the operating modes. Also, in contrast to the previous approach in [4], a modem does not need to switch off all the way to the sleeping state before switching back on again. This provides more flexibility.

As in the previous section, we assume that the workload is a Poisson process (of rate λ) of exponentially sized jobs, with mean $1/R$. We formulate the corresponding MDP to minimize the objective of a weighted sum of delay, energy due to processing, and switching costs. Decisions on switching between operating modes are made on the timescale of job arrivals and departures.

We consider a DSL modem with $c+1$ operating modes $(0, 1, 2, \ldots, c)$, with the mode c referring to full power state L0. The mode 0 refers to the sleeping state (L2 or L3, depending on which sleeping state is optimal). We also consider that each modem is associated with a finite buffer of size B. The MDP state space is $\Omega = \{0, 1, 2, \ldots, c\} \times \{0, 1, 2, \ldots, B\}$. The system is in state $i = (W, Q)$ if the modem was in mode W in the *previous* time interval and *currently* there are Q jobs in the system.

To optimize the policy that decides the sleeping and wake up process of a DSL modem, we linearly approximate, without loss of generality, the operating modes in terms of power consumption and transmit rate. Assume that P_{ON} and R correspond to the power consumption and transmit rate respectively when the modem is fully ON (mode c). When the modem is in mode W, the power consumption is equal to WP_m, and transmit rate is WR_m, where $P_m = P_{ON}/c$ and $R_m = R/c$.

The solution to the MDP is a *policy* denoted by p, indicating the optimal mode (the action) for the modem in each MDP state. The action in state $i = (W, Q)$ is the policy $p(i)$, or, equivalently, $p(W, Q)$. By the Markov structure, the policy only changes when a job arrives or departs. Note that the action is deemed to occur immediately after the transition to a given state, so that the mode in state (W, Q) is $p(W, Q)$, not W. This means that the current mode is not considered part of the state.

Similarly to the previous section, the MDP states evolve as follows:

$$W(t+\tau) = p\left(W(t), Q(t)\right) \tag{11}$$

$$Q(t+\tau) = \begin{cases} Q(t)+1 & \text{if an arrival occurs} \\ Q(t)-1 & \text{if a departure occurs,} \end{cases} \tag{12}$$

where t denotes the current time, and $t+\tau$ refers to the time of the next event. Then, the transition probabilities at state (W, Q) with policy $p(W, Q)$ and $0 < Q < B$ are:

$$Pr[(W,Q) \to (p(W,Q), Q+1)] = \frac{\lambda}{\lambda + R_m p(W,Q)}, \tag{13}$$

$$Pr[(W,Q) \to (p(W,Q), Q-1)] = \frac{R_m p(W,Q)}{\lambda + R_m p(W,Q)}. \tag{14}$$

For $Q = 0$, there are no departures, thus

$$Pr[(W,0) \to (p(W,0), 1)] = 1. \tag{15}$$

For $Q = B$, any arrivals are discarded, thus

$$Pr[(W,B) \to (p(W,B), B-1)] = 1. \tag{16}$$

Our objective is to minimize (for each MDP state) the average discounted sum of costs

$$V(i) = \min_{p(i)} \left\{ \frac{g_1(i,p(i))}{\beta + v} + g_2(i,p(i)) + \alpha \sum_{j \in \Omega} \hat{M}^{p(i)}_{i \to j} V(j) \right\}, \tag{17}$$

.

In this case we have

$$g_1(i,p(i)) = (p(i)P_{\mathrm{m}} + \mathbf{1}_{p(i)=0} P_{OFF} + rQ_i), \tag{18}$$

and

$$g_2(i,p(i)) = |W_i - p(i)|e_s + \delta_s |W_i - p(i)|\tau_s. \tag{19}$$

Note that

$$P_{OFF} = \begin{cases} P_{SL}, \text{if the optimal sleeping state is L2} \\ 0, \text{if the optimal sleeping state is L3} \end{cases} \quad (20)$$

The parameter δ_s, expressed in J/s, models the emphasis put on the time part of the switching cost relatively to the energy part. The optimal policy can be calculated by solving (17) using the value iteration algorithm [5]. Using similar arguments as in the proof of Theorem 1, we prove the following (the proof is included in Appendix):

Theorem 2: **The policy that optimally decides the transitions between different operating modes is monotone hysteretic.**

B. Simulation Results

1) Optimal Policy Structure: We illustrate the gradual transition design in an ADSL2+ downstream scenario, consisting of 6 uncoordinated lines of length 3 km deployed from a central office (CO), and 6 coordinated lines of length 1 km deployed from a remote terminal 2 km from the CO, with all subscribers collocated. The data rate of each coordinated line is 15.9 Mbps and the transmit power is 20.4dBm. The coordinated lines are initially under traffic loads shown in Table II.

TABLE II
COORDINATED LINES

Coordinated Line	Traffic Load
1	50%
2	35%
3	15%
4	10%
5	5%
6	3%

We pick those traffic loads to model underutilized DSL lines (with the exception of 50% traffic load). Initially, we consider $c = 5$ (6 operating modes), however we later study the system performance under a large range of operating modes. In order to see the structure of the optimal policy we look into line 2 (traffic load 35%), with characteristics $\frac{P_{SL}}{P_{ON}} = 0.06$ and $\frac{E^2_{ch}}{E^1_{ch}} = 1.5$. From Fig. 2 we can see that the optimal sleeping state is L2, given that for the above characteristics and a traffic load higher than 30% it is never optimal to turn off the modem to L3. We also consider a rather delay insensitive application with $r = 10$ J/job. Finally, we assume that the time required to transition between successive states is 8s, resulting in a total wake up time of 40s, and $\delta_s = 1$ J/s, thus putting equivalent emphasis on the time and energy part of the switching cost. The mean interarrival time is 10s, thus making the arrival of a job during mode transition rather not possible.

Fig. 4 demonstrates the optimal policy for line 4. We can see that the optimal ON thresholds are equal to 2, 6, 9, 13, and 16 jobs whereas the optimal OFF thresholds are equivalent to 4, 3, 2, 1, 0 jobs. This means that when the modem is OFF, it will start switching on to the first operating mode when there are 2 jobs in the buffer, then it will transition to the second

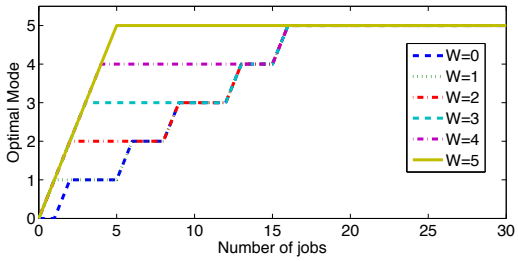


Fig. 4. Optimal policy for $c = 5$ and traffic load of 35%. The optimal sleeping state is L2. Optimal mode $W = 5$ corresponds to the modem being fully ON, and optimal mode $W = 0$ corresponds to the modem being in the lowest power mode (L2).

mode when there are 6 jobs in the buffer and so on, until there are 16 jobs in the buffer when it will actually transition to full power ON mode. Equivalently, once the modem is fully ON, it will transition to mode $W = 4$ when there are 4 jobs left in the buffer. Then it will transition to mode $W = 3$ when there are 3 jobs left in the buffer and so on, until there are 0 jobs in the buffer where it will transition to L2. Note that the traffic load of 35%, and the delay-insensitive application does not require the modem to be fully ON until there are 16 jobs in the buffer.

We should also emphasize the fact that the optimal policy does not prohibit the modem from switching on to full power from any operating mode except that it needs to go through all intermediate modes. Switching immediately to full power would be suboptimal in terms of balancing energy consumption and delay performance. For example, when the modem switches off to $W = 2$, and later on there are 11 jobs in the buffer, the modem can be designed to switch on to full power even though it is suboptimal. If the designer of the modem wishes an operation where the modem switches on directly to full power from any intermediate operating mode, this can be achieved at the ON thresholds provided by the optimal policy.

2) Optimal tradeoffs and energy savings: Here we explore the benefits that arise from increasing the granularity of control in terms of modes of operation. In particular, Fig. 5 demonstrates the Pareto optimal tradeoff between energy consumption and average delay for the coordinated line 4 (traffic load equal to 10%) in three cases; when there are 2 modes of operation available ($c = 1$), when there are 4 modes available (c=3), and when there are 6 modes available ($c = 5$). The optimal tradeoff is traced out by varying the delay sensitivity captured by parameter r. It is clear that the increased granularity of control significantly improves the Pareto optimal tradeoff, given that energy consumption reduces around 50-65% and average delay reduces from 5-65% for all delay sensitivities considered.

In addition, Fig. 6 demonstrates the average energy savings across all coordinated lines achieved by our policy when compared to the case of lines never switching off. As the number of modes increases, the energy savings increase as well, and range from 70% for 2 modes of operation ($c = 1$) to almost 90% for $c = 10$. Note, however, that a very low switching cost would result in lower energy savings, given

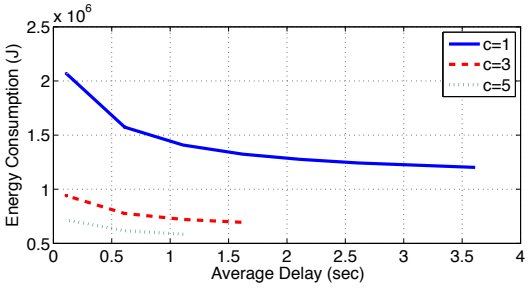


Fig. 5. Pareto optimal tradeoff between energy consumption and average delay for the coordinated line 4 under our optimal sleeping policy. Our approach results in a better energy-delay tradeoff when the number of operating modes available increases.

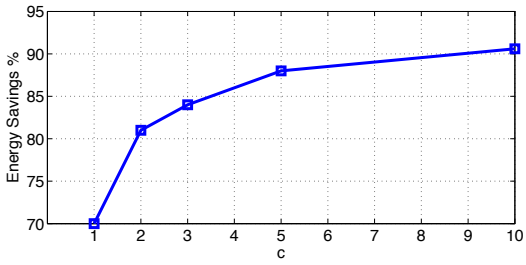


Fig. 6. Average energy savings across all coordinated lines achieved by our optimal policy for different numbers of operating modes. As the number of modes increases, the energy savings increase as well.

that the modem would have no incentive to stay OFF.

Fig.7 shows the energy savings achieved under BBF guidelines for different values of τ_{off} spent in each intermediate state. BBF guidelines perform poorly when the time spent

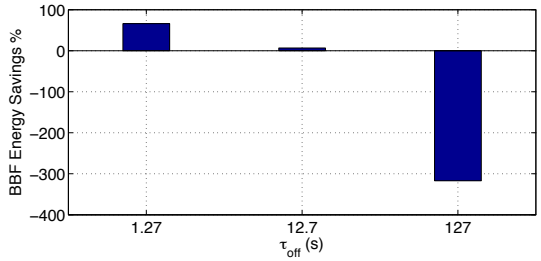


Fig. 7. Energy savings achieved by the BBF guidelines approach for different values of τ_{off}. Lower values of τ_{off} result in higher energy savings. The negative energy savings achieved for $\tau_{off} = 127$s are a result of the significant time spent in each intermediate state in addition to the switching cost incurred every time the modem switches on or off.

in each intermediate state before turning off is high ($\tau_{off} = 127$s). The performance improves as τ_{off} decreases. However, even when τ_{off} is very low ($\tau_{off} = 1.27$s), BBF guidelines perform worse than our approach with 2 modes of operation (c=1) shown in Fig. 6. The negative energy savings observed with $\tau_{off} = 127$s are not due to power consumption, which in fact decreases, but due to the extended time spent in each of the intermediate states defined by the BBF guidelines for switching off to L2.

3) Stability: Fig. 8 illustrates the improvement on maximum noise impact achieved by our optimal sleeping policy for different numbers of operating modes, in comparison with the maximum noise impact (among all frequencies) from the BBF guidelines' sleeping method. When the time spent in each

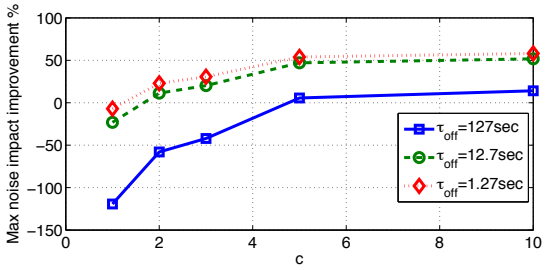


Fig. 8. Improvement on maximum noise impact achieved by our sleeping policy compared to the DSL guidelines approach for different values of τ_{off}. As τ_{off} decreases, our approach results in better stability for a large range of operating modes.

of the intermediate states defined by the BBF guidelines is relatively low ($\tau_{off} = 1.27$s or $\tau_{off} = 12.7$s), our approach performs better in terms of stability for all $c \geq 2$. In particular, when $c = 5$ or $c = 10$, our approach results in almost 50% more stable system. When $\tau_{off} = 127$s our approach performs better for all $c \geq 5$, however the amount of stability improvement is lower (around 15%). This is because the increased time spent in each of the 9 intermediate states (while switching off to L2) defined by the BBF guidelines reduces the sudden increase of electromagnetic coupling in the bundle relatively well. An approach with $\tau_{off} = 127$s as in the BBF guideline is, however, undesirable because of the poor performance in terms of energy savings.

IV. Conclusions

An optimal operation of DSL sleeping states can result in significant energy savings in DSL broadband access networks, while guaranteeing stable DSL operation. In this paper, we provide an analytic framework to optimally decide between L2 and L3 sleeping states, based on the ratio of the respective switching costs and the power consumption at L2. By using an MDP formulation, we prove the structural properties of the optimal policy that transitions a DSL line between different operating modes in order to improve the Pareto optimal tradeoff, without requiring the line to completely switch off before turning back on again. The monotone hysteretic structure of the optimal policy can result in energy savings ranging from 70% to almost 90%.

We then compare the performance of our approach with the current BBF sleeping guidelines. We show that our approach, which introduces an optimized gradual turn on procedure, along with the respective turn off procedure, can result in higher energy savings even when the BBF policy is modified to stay in each of the intermediate states 100 times shorter amount of time. In addition, our simulations present an improved stability performance of our policy compared to the BBF guidelines, even when the BBF policy operates with 127s in each intermediate state and our policy considers

only 6 operating modes. The stability improvement increases to more than 50% when our policy considers 11 operating modes ($c = 10$), which is equivalent to 9 intermediate states in BBF guidelines, and BBF policy operates with 1.27s in each intermediate state.

APPENDIX A
PROOF OF THEOREM 1

Proof: Theorem 1 of [6] proves that an optimal policy p is hysteretic if $p(\gamma, Q)$ takes the form, for some function σ and m, such as

$$p(\gamma, Q) = \text{argmin}_{\gamma' \in \mathcal{F}} \left\{ \sigma(\gamma, \gamma') + m(Q, \gamma') \right\}, \tag{21}$$

where Q is the number of customers in the server and γ is the action in $\mathcal{F}$, and if the function σ satisfies the conditions

$$\sigma(\gamma, \gamma') \leq \sigma(\gamma, c) + \sigma(c, \gamma'), \quad \forall \gamma, c, \gamma' \in \mathcal{F}, \tag{22}$$

and

$$\sigma(\gamma, \gamma) = 0, \quad \forall \gamma \in \mathcal{F}. \tag{23}$$

In our problem formulation, let $p_1(W_i^1, W_i^2, Q_i)$ denote whether the modem should be in L2 or not when we are at state i. Then for $p_2 = 0$ we can write

$$p_1(W_i^1, W_i^2, Q_i) = \text{argmin}_{\gamma \in \mathcal{F}} \{ \sigma(W_i^1, W_i^2, \gamma) + z(Q_i, \gamma) \}, \tag{24}$$

where

$$\begin{aligned} \sigma(W_i^1, W_i^2, \gamma) =& |W_i^1 - \gamma| E_{ch}^1 + W_i^2 E_{ch}^2 \\ &+ |W_i^1 - \gamma| W_i^1 \lambda \tau_1 + (W_i^2)^2 \lambda \tau_2 \end{aligned} \tag{25}$$

$$\begin{aligned} z(Q_i, \gamma) =& \frac{1}{\beta + v} (\gamma(P_{SL} - P_{ON}) + P_{ON} + rQ_i) \\ &+ \alpha \sum_{j \in \Omega} \hat{M}_{i \to j}^{\gamma} V((W_j, Q_j)). \end{aligned} \tag{26}$$

Here, $\sigma(W_i^1, W_i^2, \gamma)$ is the uniformized switching cost when the previous policy was (W_i^1, W_i^2) and the current policy is γ.

Clearly σ satisfies (22) and (23), and we can see that the properties of Theorem 1 of [6] are satisfied. Hence the optimal policy is hysteretic.

APPENDIX B
PROOF OF THEOREM 2

The hysteretic structure of the optimal policy can be proven similarly to Theorem 1. According to Theorem 3 in [6], if a policy is increasing in the queue length and operating mode and is hysteretic, then p is a monotone hysteretic policy. Hence by showing that our optimal policy is an increasing policy, we can complete the proof. According to Theorem 1 in [7], if the cost function has the following properties, then the optimal policy is increasing in queue length and operating mode.

1) The cost function is of the form $c(i, a) + s(d, a)$ where d, a are operating modes, i the number of customers in the system and $s(d, a)$ represents the switching cost from operating mode d to mode a.
2) $c(i, a)$ has the form $u(a) + hi$, where u is a function on the set of all possible operating modes.
3) The switching cost is a submodular function.
4) $s(e, a) - s(d, a) = s(e, b) - s(d, b)$, for each $a < b \leq d < e$ in the set of operating modes.

In our problem we can write the cost function as

$$g(i, p(i)) = u(p(i)) + rQ_i + s(W_i, p(i)), \tag{27}$$

where

$$\begin{aligned} u(p_1(i)) &= p(i) P_{\text{m}} + \mathbf{1}_{p(i)=0} P_{OFF} \\ s(W_i, p(i)) &= |W_i - p(i)| e_s + \delta_s |W_i - p(i)| \tau_s. \end{aligned}$$

Also, the switching cost $s(W_i, p(i))$ is submodular. For each $a < b \leq d < e$,

$$\begin{aligned} s(e, a) - s(d, a) &= |e - a| e_s - |d - a| e_s \\ &\quad + \delta |e - a| \tau_s - \delta |d - a| \tau_s \\ &= (e - d) e_s + \delta (e - d) \tau_s. \end{aligned} \tag{28}$$

$$\begin{aligned} s(e, b) - s(d, b) &= |e - b| e_s - |d - b| e_s \\ &\quad + \delta |e - b| \tau_s - \delta |d - b| \tau_s \\ &= (e - d) e_s + \delta (e - d) \tau_s. \end{aligned} \tag{29}$$

Hence,

$$s(e, a) - s(d, a) = s(e, b) - s(d, b). \tag{30}$$

In order for s to be submodular we need to have

$$\begin{aligned} s(\max(a, c), \max(b, d)) + s(\min(a, c), \min(b, d)) \\ \leq s(a, b) + s(c, d) \end{aligned} \tag{31}$$

for two points (a, b) and (c, d).

By checking all possible combinations, we verify that the above submodularity condition is indeed satisfied.

From all of the above, we conclude that the optimal hysteretic policy p is increasing. Hence the optimal policy is monotone hysteretic.

■

REFERENCES

[1] J. M. Cioffi, S. Jagannathan, W. Lee, H. Zou, A. Chowdhery, W. Rhee, G. Ginis, P. Silverman, "Greener Copper with Dynamic Spectrum Management", in *AccessNets*, Las Vegas, NV, USA, October 2008.

[2] P. Tsiaflakis, Y. Yi, M. Chiang, and M. Moonen,"Green DSL: Energy-efficient DSM", in *Proc. of IEEE Internat. Conf. on Communications (ICC 2009)*, Dresden, Germany, June 2009.

[3] G. Ginis, "Low-Power Modes for ADSL2 and ADSL2+", White paper SPAA021, Broadband Communications Group, Texas Instruments, January 2005.

[4] I. Kamitsos, P. Tsiaflakis, S. Ha, and M. Chiang, "Stable Sleeping in DSL Broadband Access: Feasibility and Tradeoffs", in *Proc. IEEE Globecom*, Houston, TX, December 2011.

[5] D.P. Bertsekas, "Dynamic Programming and Optimal Control", vol. 2, Athena Scientific, Belmont, MA, USA, 2007.

[6] S.K. Hipp, and U.D. Holzbaur, "Decision Processes with Monotone Hysteretic Policie", *Operations Research* vol. 36, no. 4, pp. 585-588, July-August 1988.

[7] F.V. Lu and R.F. Serfozo, "M/M/1 Queueing Decision Processes with Monotone Hysteretic Optimal Policies", *Operations Research*, vol. 32, no. 5, September-October 1984.

[8] Broadband Forum TR-202, "ADSL2/ADSL2plus Low-Power Mode Guidelines", February 2010.

[9] M.Wolkerstorfer, D. Statovci, and T. Nordstrom, "Energy-saving by low-power modes in ADSL2", *Computer Networks*, vol. 56, no. 10, pp 2468-2480, July 2012.