Stable Sleeping in DSL Broadband Access: Feasibility and Tradeoffs

Ioannis Kamitsos[1],Paschalis Tsiaflakis[2], Sangtae Ha[1] and Mung Chiang[1]
[1]Princeton University, USA, [2]CORE, Université Catholique de Louvain, Belgium

***Abstract*—Energy efficient and stable operation of the DSL broadband access infrastructure has become an essential part of the emerging trend towards green communications. One promising means to obtain energy savings is the use of low power "sleep modes", putting DSL modems to sleep when they are not used. Executing the optimal sleeping policies is, however, not straightforward since turning the modem ON and OFF consumes both energy and time, and it also impacts the stability of the DSL network. We present an analytic framework providing optimal sleeping policies that achieve a Pareto-optimal tradeoff between energy consumption and delay performance. Furthermore, we present mechanisms achieving stable sleep mode operation that improve overall stability of DSL systems. Using a realistic DSL simulator, we demonstrate the three-way tradeoff between energy consumption, delay performance and stability.**

I. Introduction

Digital subscriber line (DSL) technology is currently the most widely used wireline Internet access technology, with a global market share of more than 64%, corresponding to more than 300 million subscribers. Given this global penetration and the typical always-ON strategy of off-the-shelf DSL modems, reducing energy consumption of the DSL broadband access network has become an essential part of the emerging trend towards Green Information and Communication Technology (Green ICT), such as the 50% energy reduction targeted by EU on broadband access equipments. One approach to reduce the energy consumption of DSL access networks, is to reduce the always-ON transmit powers using energy efficient dynamic spectrum management (DSM) [1][2][3], also referred to as Green DSL. It jointly optimizes the transmit power of each DSL modem in order to minimize the overall transmit power consumption, subject to certain quality of service requirements such as minimum data rates and fairness objectives.

In this paper, we focus on an alternative approach: the optimization of sleep mode activation, in which turning on and off policies are optimally decided based on the arrival data load. This method can be highly effective because a significant amount of energy is wasted due to low utilization in light of bursty IP traffic; for example, a typical DSLAM is on average not transmitting data during 80% of time [5]. However, the design of an optimal sleeping policy is nontrivial for two reasons: (1) the best trade-off point depends on the relative importance of energy-saving versus delay performance, and (2) it takes both time and energy to turn a modem on or off. The latter implies that the optimal policy may be to remain ON but idle rather than turning off, if there is a high chance of having to turn on again soon.

This work was in part sponsored by NSF NeTS and AFOSR MURI grants.

In addition, an important concern, and a bottleneck issue in practice, of switching DSL modems on and off is the generation of *time-varying* crosstalk in the DSL access network [6], because of the presence of electromagnetic coupling, i.e. crosstalk, among different lines within the same cable bundle. This may cause network instability which should be avoided as much as possible. A solution for sleep mode optimization should thus also consider the *stability* issue, in order to obtain a desirable three-way trade-off among energy consumption, delay performance and stability of the network.

In this paper we develop a novel analytic framework that computes optimal ON-OFF policies for DSL modems, that achieve a Pareto-optimal tradeoff between energy consumption and delay performance by exploiting sleep modes available in DSL standards. We pay close attention to the impact on stability, where we explore solutions along spatial and temporal dimension that can significantly improve the stability of sleep mode policies. We evaluate and analyze potential energy savings with the optimal sleep policies. Numerical results show that energy savings of about 85% can be achieved for the traffic load of 10%.

II. DSL system model

A standard DSL system model consists of a cable bundle of $M+N$ DSL lines (modems), in which lines $n \in \mathcal{N} = \{1,\dots,N\}$ are coordinated by a spectrum management center (SMC), whereas lines $m \in \mathcal{M} = \{N+1,\dots,N+M\}$ are not coordinated (e.g., they are bundled together while they belong to different service providers). We assume standard Discrete Multi-Tone (DMT) modulation for the DSL modems with perfect synchronization. No joint signal coordination is assumed between the transmitting and receiving DSL modems. The crosstalk interference between DSL modems is treated as additive white Gaussian noise, resulting in the following achievable data rate expression for each modem n:

$$R^n = \sum_{k=1}^{K} \log_2\left(1 + \frac{1}{\Gamma_{co}\Gamma^n}\frac{h_k^{n,n} w_k^n s_k^n}{\mathrm{int}_k^n}\right), \qquad (1)$$

with $\mathrm{int}_k^n = \sum_{i\neq n} h_k^{n,i} w_k^i s_k^i + \sigma_k^n$ denoting the received interference, $h_k^{n,n}$ the channel gain of line n on tone k, $h_k^{n,i}$ the crosstalk channel gain from line i to line n on tone k, s_k^i the transmit power of line i on tone k, σ_k^i the noise power of line i on tone k, Γ_{co} the SNR gap of the code, Γ^n the noise margin of line n, and K the number of tones.

We also include *power management mode parameters* w_k^i for each line i on tone k. DSL standards namely define a number of power management modes for DSL modems (lines), known as L0, L2 and L3 power management modes. For instance, L0 mode corresponds to the DSL line being fully functional, whereas L3 mode corresponds to the case where no signal is transmitted on the line. For this latter mode, significant power savings can be obtained, depending on whether the corresponding chipsets (or ports of the chipset) are placed in a 'monitor' state or on a 'disabled' state, where the latter provides larger savings [6]. We make an abstraction of the power management modes whose implementations are typically different among chipsets. We model them by the parameters w_k^i, which can only be coordinated for the coordinated lines $i \in \mathcal{N}$, i.e., fixed settings for the uncoordinated lines $w_k^i = 1, \forall k, \forall i \in \mathcal{M}$. The ON mode of a coordinated DSL modem n is defined as $w_k^n = 1, \forall k$, whereas the OFF or sleep mode is defined as $w_k^n = 0, \forall k$. We present an extension towards intermediate power management modes in Section V. Note that the ON-OFF power management modes of the N coordinated DSL modems can be individually controlled depending on their particular data load.

The data rates and corresponding power consumptions of the coordinated DSL modems depend on the chosen transmit powers and the power management modes, and are denoted as R^n and P^n_{ON} for the ON mode of modem n, respectively. The data rate of the OFF mode is zero, and the corresponding power consumption is also chosen to be zero, although any other number smaller than P^n_{ON} can also be allowed. As we focus on power management mode coordination, i.e., sleeping coordination, we assume that the transmit spectra are fixed, and for stability they respect a noise margin Γ^n when all lines are ON.

The introduction of ON-OFF sleeping mode coordination results in time-varying crosstalk, which can significantly impact the stability of the M uncoordinated lines. A proper definition of stability is however not straightforward and depends on many DSL modem implementation issues. Generally, if a line m undergoes an increase of crosstalk interference that is larger than the noise margin Γ^m, then the line can get unstable. However, practical rescue mechanisms such as bit swapping (BS) and seamless rate adaptation (SRA) can mitigate instability up to a certain degree, depending on their reaction speed, the target rates, and other practicalities. Therefore, we define the noise impact z_k^m (for tone k and line m) as *stability measure*, which represents the impact on the uncoordinated lines caused by the OFF-ON transitions of coordinated lines:

$$z_k^m = \frac{(\text{int}_k^m)_{\text{after transition}}}{(\text{int}_k^m)_{\text{before transition}}}, m \in \mathcal{M}, k \ldots \{1, \ldots, K\}, \quad (2)$$

with the nominator and denominator referring to the interference of the uncoordinated lines after and before OFF $\rightarrow$ ON transitions, respectively. Lower noise impact results in a more stable line. Between successive transitions, uncoordinated lines can adapt to the new interference level by using the practical rescue mechanisms mentioned earlier in order to recover their noise margin. Definition (2) allows to define stability improving solutions resulting in stable sleep operation as will be discussed in Section V.

III. Analytic framework for ON-OFF power mode optimization

We assume that each of the coordinated DSL modems has a finite buffer of size B and two modes, ON (working) and OFF (sleeping), where $W = ON$ corresponds to $w_k^n = 1, \forall k$, and $W = OFF$ corresponds to $w_k^n = 0, \forall k$. Note that for clarity we will continue with W instead of the power management mode parameters w_k^n. Furthermore we leave out the modem index n for notational simplicity, as this section focuses on modeling ON-OFF behavior and its corresponding energy-delay trade-off for one modem.

We assume that jobs arrive to the DSL modem according to a Poisson process with rate λ; The jobs have independent, identically exponentially distributed service time with mean $1/R$. The memoryless property of the arrivals allow us to formulate the problem as a continuous time Markov Decision Process (MDP) problem. The notation that will be used in the problem is summarized in Table 1.

The state space for each modem is $\Omega = \{ON, OFF\} \times \{0, 1, 2, \ldots, B\}$. The modem is in state $i = (W, Q)$ if there are Q jobs in the buffer and the policy was W *before* the most recent arrival or departure. We often write $i(t)$, $W(t)$ and $Q(t)$ to denote the state at continuous time t. The action space $\mathcal{F}$ includes only two actions, i.e., $\mathcal{F} = \{ON, OFF\}$. The rate of doing work depends only on the action: If the action is ON, then the modem is in the ON mode, and jobs are processed at rate R; otherwise, the modem is OFF and no jobs are processed. When the modem is ON, it consumes a constant power P_{ON}, and when it is OFF, it consumes no power. Turning the modem from OFF to ON or from ON to OFF each consumes energy E_{ch}. Also, turning a modem ON takes time, which is denoted by τ. For instance, $\tau = 20sec$ or $\tau = 3sec$ when fast initialization is used in ADSL2+ [6].

As in [7], our objective is to minimize a weighted sum of the expected response time and the energy consumed per job, while at the same time, whenever there is an ON $\rightarrow$ OFF or OFF $\rightarrow$ ON switching an energy cost of E_{ch} is considered. However, in contrast to our previous work on energy reduction in servers [7], here we take into account the time it takes for the modem to wake up. When we put the line to sleep for a long time we gain in energy saving but we lose in delay and vice versa. In addition, stability of sleeping in a crosstalk environment is a new dimension in this work.

The solution to the MDP is a (sleeping) *policy* denoted by p, saying which action to perform (which mode to enter) in each state. In other words, the objective is to find an optimal action for each state. The action, either ON or OFF, in state $i = (W, Q)$ is denoted by the binary value $p(i)$, or, equivalently, $p(W, Q)$. By the Markovian structure, the policy only changes when a job arrives or departs.

The states evolve as $W(t+1) = p\,(W(t), Q(t))$. If an arrival occurs then $Q(t+1) = Q(t) + 1$. However, if a departure occurs then $Q(t+1) = Q(t) - 1$. The transition probabilities

TABLE I
MAIN NOTATION

Symbol	Meaning [Units]
P_{ON}	Power consumed when DSL modem is ON [Watts]
B	Buffer size [Number of jobs]
$i = (W, Q)$	i-th state in the state space
W	Defines whether the DSL modem is ON or OFF
Q	Queue length [Number of jobs]
R	Service rate [jobs/sec]
λ	Arrival rate [jobs/sec]
$p(i)$	Policy for state i
E_{ch}	Energy consumed on transitions ON→ OFF and OFF→ ON [kWh]
τ	Time required for DSL modem to wake up
$M_{i\to j}^{p(i)}$	Transition probability of moving from state i to state j under policy $p(i)$
$V(i)$	Average cost of state i
$g(i, p(i))$	Cost per transition of state i under policy $p(i)$
δ	Scaling parameter [Joules/job]
r	Congestion cost coefficient [W/job]
w_k^n	Power management mode parameter for DSL modem n on tone k
α	Discount factor

at state (W, Q) with policy $p(W, Q)$ and $0 < Q < B$ are

$$Pr[(W, Q) \to (p(W, Q), Q + 1)] = \frac{\lambda}{\lambda + Rp(W, Q)}, \quad (3)$$

$$Pr[(W, Q) \to (p(W, Q), Q - 1)] = \frac{Rp(W, Q)}{\lambda + Rp(W, Q)}. \quad (4)$$

For $Q = 0$, there are no departures, hence $Pr[(W, 0) \to (p(W, 0), 1)] = 1$.

For $Q = B$, any arrivals are discarded, hence $Pr[(W, B) \to (p(W, B), B - 1)] = 1$.

We define $g(i, p(i))$ as the running cost per unit time incurred at state i under policy $p(i)$. It consists of two components; the first, $g_1(i, p(i))$, is a weighted sum of the energy cost due to the DSL modem being ON and the delay cost incurred by the jobs waiting to be processed. Hence,

$$g_1(i, p(i)) = p(i)P_{\text{ON}} + rQ_i, \quad (5)$$

where r is the tradeoff parameter between the energy cost and the cost of delay. The second component, $g_2(i, p(i))$, captures the switching cost. Switching cost is the energy that is spent for OFF/ON switchings. We also take into account the time required for the modem to wake up. The latter is expressed in terms of $\lambda\tau$ jobs that accumulate in the buffer while the modem wakes up. Hence,

$$g_2(i, p(i)) = |W_i - p(i)|E_{\text{ch}} + \delta|p(i) - W_i|(1 - W_i)\lambda\tau, \quad (6)$$

where δ is the scaling parameter. Note that here we take $W = 1$ for ON mode and $W = 0$ for OFF mode.

The energy is saved by reducing the number of times we expend the energy E_{ch} required to turn ON/OFF the modem. This energy/delay tradeoff is captured by the congestion cost slope coefficient r, which measures how much emphasis is put on the congestion cost. For example, if we have an application that is rather delay insensitive, then r is relatively small. This allows the modem to sleep longer and decreases the cost incurred by the line being in ON mode. As a result, the congestion/delay cost is increased.

The problem thus far is a continuous time problem where the transition rates are time-varying. By applying the standard technique of "uniformization" ([7], [8]), our model can be analyzed in the discrete time domain, making it more analytically tractable. The uniformized cost per transition is given by

$$\hat{g}(i, p(i)) = \frac{1}{\beta + v}(p(i)P_{\text{ON}} + rQ_i) \quad (7)$$
$$+|W_i - p(i)|E_{\text{ch}} + \delta|p(i) - W_i|(1 - W_i)\lambda\tau$$

where v is the uniform transition rate and is given by $v = \lambda + R$.

Our objective is then to minimize the average discounted sum of costs, given by

$$V(i) = \min_{p(i)}\{\frac{g_1(i, p(i))}{\beta + v} + g_2(i, p(i)) + \alpha \sum_{j\in\Omega} \hat{M}_{i\to j}^{p(i)} V(j)\} \quad (8)$$

Here $\hat{M}_{i\to j}^{p(i)}$ denotes the "uniform" transition probability going from state i to state j under policy $p(i)$ [8].

The optimal policy can be calculated by solving (8) with the value iteration algorithm [8].

IV. Optimal Policy's Structure

The following definitions are useful for the analytic results presented in this section.

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

Definition 2: [9] 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 service levels, 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

$$p(i, d) = d, \text{if } l_d \le i < u_{d+1}, d = 1, \ldots, |A|$$
$$p(i, d) = p(i, d + 1), \text{if } i \ge u_{d+1}, d = 1, \ldots, |A|$$
$$p(i, d) = p(i, d - 1), \text{if } i < l_d, d = 2, \ldots, |A|$$

Theorem 1: The optimal policy for the above MDP is a monotone hysteretic policy.

The proof of the above theorem is provided in Appendix I. Given that there are two modes of operation, the monotone hysteretic sleeping policy is characterized by two queue length thresholds, one for turning the modem on and one for turning it off. Such a policy is useful, since it is particularly easy to implement in a practical system.

V. Stable sleeping

The introduction of sleeping modes for coordinated DSL lines results in an increased noise impact on the uncoordinated lines. The concern about sleeping mode induced instability has long been a bottleneck issue. To improve the stability behavior of sleeping mode coordination, we propose the following solutions along two different dimensions:

1) **Spatial dimension**: The number of coordinated lines that can wake up simultaneously is reduced from N to a *scalar* number between $N-1$ and 1.
2) **Temporal dimension**: The wake-up of a DSL modem from OFF to ON occurs gradually following some increasing time *function* instead of being a step function, and is implemented by discrete increasing intermediate power management states with fixed time intervals equal to the OFF-ON wake up time, e.g., $w_k^n \in \{0, 1/8, 1/4, 1/2, 1\}, \forall k, \forall n \in \mathcal{N}$.

These dimensions allow to reduce the noise impact (2) on the uncoordinated lines and thus improve the stability of these lines. In Section VI we demonstrate an improved *stable sleep mode operation* by exploiting the time and spatial dimensions.

The following intuitive result (proven in Appendix II) particularly shows how the time dimension solution improves the stability of the uncoordinated lines.

Theorem 2: Suppose that for coordinated line n there are i intermediate states between modes OFF→ON characterized by transmit powers $\{0, \frac{1}{2^i}, \frac{1}{2^{i-1}}, \ldots, \frac{1}{2}, 1\}P_{\text{ON}}$. Then, as the number of intermediate states i increases, the noise impact on each uncoordinated line m decreases.

VI. Numerical results

Practical DSL setup: For the numerical evaluation of our sleeping policies we consider an upstream DSL setup that consists of 4 DSL lines; line 1 is 1200m, line 2 is 1000m, line 3 is 800m and line 4 is 600m. Lines 3 and 4 are under our control and can be coordinated (they belong to the same ISP), however lines 1 and 2 are uncoordinated. The transmit spectra of the uncoordinated lines are fixed at 60 dBm/Hz. The transmit spectra of the coordinated lines are optimized for energy efficiency using the following energy efficient DSM formulation [3], which is solved with the DSB spectrum optimization algorithm [4]:

$$\begin{array}{ll} \underset{\text{transmit powers}}{\text{maximize}} & \sum_{n \in \mathcal{N}} R^n \\ \text{subject to} & P^n \leq 0.5 P^{n,\text{tot}}, \quad \forall n, \end{array} \tag{9}$$

In the above formulation, P^n and $P^{n,\text{tot}}$ correspond to the per modem transmit power and available power budget defined by DSL standards, respectively. This results in both coordinated lines 3 and 4 having a transmit power consumption of $0.5P^{n,\text{tot}}$. Furthermore, line 4 has a transmit rate 2.5 times higher than the transmit rate of line 3, i.e $R^4 = 2.5R^3$. The traffic load for line 3 is assumed to be 50% and for line 4 is 20% (under the same arrival rate).

A. Optimal ON-OFF policy

For the first set of simulations we consider $r = 10$ W/job, corresponding to a rather delay tolerant application. The scaling parameter was chosen to be $\delta = 1000$ J/job. We will later demonstrate how the scaling parameter affects the structure of the optimal sleeping policy. We also assumed a wake up process characterized by a switching time of approximately $\tau = 20$ sec [6]. Fig. 1 demonstrates the optimal policy of line 4. We can observe that the monotone hysteretic structure predicted by Theorem 1 is confirmed. The modem of line 4 sleeps until 8 jobs are accumulated in the buffer and then wakes up. Once awake, the modem sleeps until all the jobs in the buffer are processed. The optimal policy of line 3 has the same structure but the ON threshold is 6 jobs. Given the higher traffic load, line 3 wakes up sooner to process the jobs that are accumulated in the buffer. Fig. 2 shows how

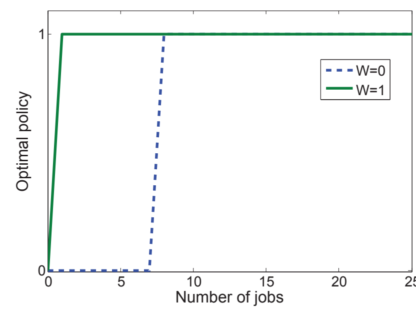

Fig. 1. Optimal sleeping policy for modem of line 4. The modem sleeps until 8 jobs are accumulated in the buffer.

the optimal ON thresholds of lines 3 and 4 vary with the scaling parameter δ. As the scaling parameter increases, the ON threshold increases for both lines. Increasing the scaling parameter makes the switching cost higher, hence modems stay in sleeping mode for longer periods in order to avoid too frequent ON/OFF switchings. Also, for any value of δ, the optimal ON threshold of line 3 is lower than that of line 4, since line 3 is characterized by a higher traffic load. Compared

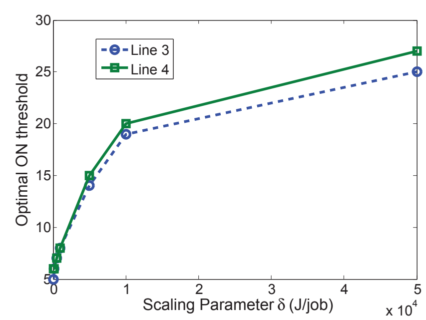

Fig. 2. Optimal ON threshold for different values of δ. The higher the scaling parameter, the higher the ON threshold for both lines, since both lines try to avoid too frequent ON/OFF transitions for higher switching costs.

to our previous work in [7], the consideration of the time it takes a modem to wake up increases the hysteresis of the optimal policy, given that the switching cost becomes higher.

B. Optimal Trade-off: Energy vs Delay vs Stability

As a key practical consideration, sleeping policies can cause stability issues on the uncoordinated lines 1 and 2 since the sudden wake up of lines 3 and 4 produces time varying crosstalk. We will demonstrate the improvement of the stability of lines 1 and 2 by exploiting the temporal and spatial dimension (as discussed in Section V), where the latter corresponds to reducing the number of lines that can wake up simultaneously to 1. Regarding the time domain, the modem's wake up process is made gradual, via an exponential waking up function. Specifically, when there are i intermediate states between OFF→ON, the transmit powers are $\{0, \frac{1}{2^i}, \frac{1}{2^{i-1}}, \ldots, \frac{1}{2}, 1\}P_{\text{ON}}$. Note that the transmit spectra of the intermediate states correspond to simple power back-off versions of the optimal transmit spectra obtained from (9). As noted before, we assume a large enough fixed time interval between consequent waking up states that allows the uncoordinated lines to adapt to the new interference level.

Fig. 3 shows the Pareto optimal energy consumption - average delay tradeoff of line 3 when the monotone hysteretic sleeping policy is considered. The tradeoff graph is derived by varying the tradeoff parameter r. Each value of this parameter corresponds to a different optimal policy. It is clear that the sleeping policy results in energy being saved in cost of increased delay. Note that the more intermediate states considered, the worse the system performance is, since for given energy consumption the average delay increases. This is because of the fact that the more intermediate states we consider, the more time it takes for the modem to wake up to full ON power (for example, given our assumptions, one intermediate state doubles the wake up time), hence the increase in delay. The increase in delay is the price we need to

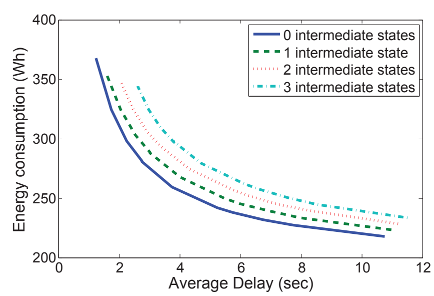

Fig. 3. Pareto optimal tradeoff between energy consumption and average delay of line 3. The more intermediate states introduced, the higher the average delay for given energy consumption level.

pay in order to reduce the noise impact on the uncoordinated lines 1 and 2, and thus to improve the stability with respect to the direct OFF → ON case.

Fig. 4 shows the improvement of the time averaged maximum noise impact on both lines 1 and 2, from the introduction of the intermediate states in the wake up process, compared to the case where there is a direct OFF → ON switching. The improvement is significant and increases with the number of intermediate states. Specifically, it is around 35% for 1 intermediate state, around 60% for 2 states and around 70% for 3 states. The additional introduction of the spatial dimension constraint that lines 3 and 4 cannot wake up at the same time improves the stability even further for any number of intermediate states. Fig. 3 and 4 demonstrate the trade-off between energy consumption, delay performance and stability. Finally, Fig 5 shows how the energy savings of line 3, achieved by our sleeping policy, vary with traffic load. These are compared to the energy savings of a baseline policy that wakes up the modem as soon as there is one job to be processed and puts the modem to sleep when there are no jobs in the

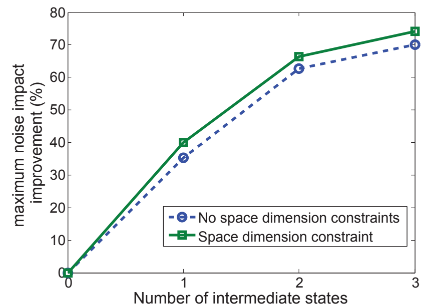

Fig. 4. Improvement of the time averaged maximum noise impact on the uncoordinated lines by the introduction of the intermediate states and the spatial dimension constraint. The more the intermediate states, the higher the improvement. The latter increases more when lines are not allowed to wake up at the same time.

buffer. The energy savings are expressed in percentage of the energy that would be consumed if the modem was always ON. For the optimal policy and for low traffic load the energy savings are significant (85% for traffic load of 10%), given that that the modem does not have to be ON for a long time to process incoming jobs. However, as the traffic load increases, the energy savings naturally decrease. Also, for any traffic load, the optimal policy results in higher energy savings than the baseline policy. It is notable that, for some traffic loads, the baseline policy demonstrates "negative" energy savings, implying that keeping the modem ON would consume less energy than having an ON threshold of one job, which results in too frequent ON/OFF switchings.

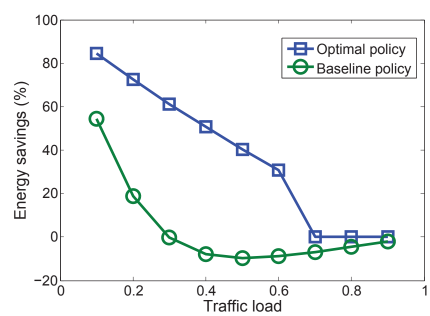

Fig. 5. Energy savings of line 3 over traffic load.As the traffic load increases, the modem has to stay ON for longer time, thus decreasing the energy savings. The optimal policy results in higher energy savings for all traffic loads compared to a baseline policy that turns the modem ON as soon as there is 1 job in the buffer.

VII. Conclusion

The use of DSL sleeping modes offers a great potential for decreasing energy consumption of the DSL broadband access infrastructure. In this work, we develop optimal and stable sleeping mode policies. The problem formulation is based on a Markov Decision Process (MDP) with easily tunable parameters that allow to properly model DSL sleeping modes. Our MDP-based optimal policies result in a Pareto-optimal tradeoff between energy consumption and delay performance. It is shown that the optimal policy has a monotone hysteretic structure, characterized by two thresholds. Numerical results demonstrate energy savings of up to 85% for a traffic load of 10%.

Stability improving mechanisms are then proposed by considering the impact on uncoordinated lines. These consist of (1) reducing the number of DSL modems that can simultaneously be turned on (spatial dimension), and (2) smoothing the wake-up process of DSL modems in time by introducing intermediate power states (temporal dimension). The benefits of these mechanisms are evaluated by numerical simulation results, demonstrating a three-way tradeoff between energy consumption, delay performance and stability. More specifically, the introduction of intermediate states in OFF $\rightarrow$ ON transitions significantly improves stability, given that the time averaged maximum noise impact on the uncoordinated lines improves by a factor up to 75% when considering 3 intermediate states.

Appendix I: Proof of Theorem 1

Proof: First we will show that the optimal policy is hysteretic. Specifically, Theorem 1 of [9] proves that the optimal policy is hysteretic if $p(\gamma, Q)$ takes the form, for some function σ and m, such as

$$p(\gamma, Q) = \text{argmin}_{\gamma' \in \mathcal{F}} \left\{\sigma(\gamma, \gamma') + m(Q, \gamma')\right\}, \tag{10}$$

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

$$\sigma(\gamma, \gamma') \leq \sigma(\gamma, c) + \sigma(c, \gamma'), \quad \forall \gamma, c, \gamma' \in \mathcal{F}, \tag{11}$$

and

$$\sigma(\gamma, \gamma) = 0, \quad \forall \gamma \in \mathcal{F}. \tag{12}$$

In our problem formulation, if $p^*(W_i, Q_i)$ is the optimal policy of state i, then we can write

$$p^*(W_i, Q_i) = \text{argmin}_{\gamma \in \mathcal{F}} \{\sigma(W_i, \gamma) + m(Q_i, \gamma)\}, \tag{13}$$

where

$$\sigma(W_i, \gamma) = |W_i - \gamma| E_{\text{ch}} + |\gamma - W_i|(1 - W_i)\lambda\tau \tag{14}$$

$$m(Q_i, \gamma) = \frac{1}{\beta + v}(\gamma P_{\text{ON}} + rQ_i) + \alpha \sum_{j \in \Omega} \hat{M}^{\gamma}_{i \rightarrow j} V((W_j, Q_j)). \tag{15}$$

Clearly σ satisfies (11) and (12), and we can see that the properties of Theorem 1 of [9] are satisfied. Hence the optimal policy is hysteretic.

According to Theorem 3 in [9], if a policy is increasing in the queue length and service level and hysteretic, then p is a monotone hysteretic policy. Hence, what remains to show is that our optimal policy is increasing. According to Theorem 1 in [10], if the cost function has the following properties then the optimal policy is increasing in queue length and service level.

1) The cost function is of the form $c(i, a) + s(d, a)$ where d, a are service levels, i the number of jobs in the queue and $s(d, a)$ represents the switching cost from service level d to level a.
2) $c(i, a)$ has the form $u(a) + hi$, where u is a function on the set of all possible service levels.
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 service levels.

In our problem we can write the cost per transition as

$$g(i, p(i)) = u(p(i)) + hi + s(W_i, p(i)), \tag{16}$$

where $u(p(i)) = p(i)P_{\text{ON}}$ and $s(W_i, p(i)) = |p(i) - W_i|E_{\text{ch}} + |p(i) - W_i|(1 - W_i)\lambda\tau$ Also, the switching cost $s(W_i, p(i))$ is submodular and for each $a < b \leq d < e$ in the set of service levels we can see that $s(e, a) - s(d, a) = s(e, b) - s(d, b)$ From the above we can conclude that the optimal hysteretic policy p is increasing. Hence the optimal policy is monotone hysteretic. ∎

Appendix II: Proof of Theorem 2

Proof: Suppose that coordinated line n is at a state where the transmit power is $w^n_k s^n_k$, for tone k and that it transits to the next state where the transmit power is $(w^n_k + \epsilon)s^n_k$. Then, from (1) and (2) we can see that the noise impact of this transition on the uncoordinated line m is of the form

$$z^m_k = 1 + \frac{\epsilon}{a + w^n_k}, \tag{17}$$

where a is a positive constant. If the number of intermediate states increase from i to $i + 1$, this replaces the transition $0 \rightarrow \frac{1}{2^i}$ (transition A) with the transitions $0 \rightarrow \frac{1}{2^{i+1}}$ (transition B) and $\frac{1}{2^{i+1}} \rightarrow \frac{1}{2^i}$ (transition C). From (17), transition A has a noise impact of $1 + \frac{1}{a2^i}$. Similarly, transitions B and C have a noise impact of $1 + \frac{1}{a2^{i+1}}$ and $1 + \frac{1}{a2^{i+1}+1}$ respectively. Both of these transitions have a lower noise impact than transition A. Hence,

$$z^m_{k,i+1} \leq z^m_{k,i}$$

where $z^m_{k,i}$ is the noise impact on line m and tone k when there are i intermediate states. ∎

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] P. Tsiaflakis, Y. Yi, M. Chiang, and M. Moonen, "Fair greening for DSL broadband access", in *ACM Sigmetrics/Performance GreenMetrics Workshop*, Seattle, WA, USA, June 2009.

[4] P. Tsiaflakis, M. Diehl, M. Moonen, "Distributed Spectrum Management Algorithms for Multiuser DSL Networks", in IEEE Trans. on Signal Processing, vol. 56, no. 10, pp. 48254843, Oct 2008.

[5] ETSI, "Environmental Engineering; The reduction of energy consumption in telecommunications equipment and related infrastructure,", ETSI TR 102 530 v1.1.1, June 2008.

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

[7] I. Kamitsos, L. Andrew, H. Kim, and M. Chiang, "Optimal Sleeping Patterns for Serving Delay Tolerant Jobs", in *Proc. 1st International Conference on Energy Aware Computing and Networking, e-Energy*, pp. 31-40, Passau, Germany, April 2010.

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

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

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