Stable Sleep Mode Optimization for Energy Efficient DSL

Ioannis Kamitsos, Paschalis Tsiaflakis, *Member, IEEE*, Ken J. Kerpez, *Fellow, IEEE*,
Sangtae Ha, *Senior Member, IEEE*, and Mung Chiang, *Fellow, IEEE*

***Abstract*—In this paper, we optimize the use of existing DSL low-power sleeping modes (L2 and L3) in order to improve the energy efficiency of DSL access networks. Given that switching ON a DSL line can cause instability by increasing the amount of time-varying crosstalk in the cable bundle, and that it takes energy and time to switch a DSL line ON and OFF, we develop a method to optimally choose the appropriate sleeping state based on the line and traffic characteristics. We further develop and prove the structural properties of the optimal policy for switching to the appropriate sleeping state, allowing transitions between submodes with different power and transmit rate characteristics. We also present techniques that guarantee stable sleep mode operation. Using a realistic DSL simulator, we demonstrate the three-way tradeoff among energy consumption, delay performance, and stability. The increased flexibility of control introduced by our approach improves the energy-delay Pareto optimal tradeoff, and results in a more energy efficient and stable DSL operation compared to existing power saving policies.**

***Index Terms*—Broadband, DSL, energy efficiency, green.**

I. INTRODUCTION

DSL IS THE predominant wireline Internet access technology with more than 350 million lines worldwide [1]. In this context, saving energy in DSL infrastructure is essential in the overall movement to "green" the access network. Exploiting DSL low-power modes can be highly effective because a significant amount of energy is wasted due to low utilization; for example, a typical DSLAM is fully on but does not transmit data 80% of time [2] and yet most DSL lines do not switch off. The potential is further enhanced due to the increasing burstiness of broadband traffic and delay tolerance of cloud applications.

A simple "sleeping" policy would wake up a DSL line as soon as there are data to be transmitted and then put it to sleep when the transmission is over. Designing an optimal "sleeping" policy is complex for three reasons: (1) the best trade-off point depends on the relative importance of energy-saving versus delay performance, (2) it takes both time and energy to switch a line on or off, and (3) DSL lines sharing a cable binder crosstalk into each other, so switching DSL lines on and off generates *time-varying* crosstalk [3] which can cause errors and instability since the lines receiving the crosstalk adapt at start-up assuming stationary noise. In this work, stability is related to time varying noise impact through crosstalk coupling. More specifically, the larger the noise impact magnitude and variation, the more symbol errors occur, and as a result, the less stable DSL lines become as they will go out of showtime and re-initialize every time a threshold value of symbol errors is exceeded. Designing sleeping policies should thus take into account the *stability* metric, aiming at achieving a three-way tradeoff between energy consumption, delay performance and stability.

Part of the material in this manuscript has been published in [4] and [5]. Extending those works, this paper submission presents a unified theory of stable sleeping in DSL infrastructure. While each one of our conference papers focused on specific aspects of DSL stable sleeping, this work presents a comprehensive and thorough theoretical analysis of optimizing the tradeoff between energy consumption, delay, and stability; and documents the reasons behind the development of the ITU-T VDSL Low Power Mode (LPM) [6].

The contributions of this paper are the following:

1) We initially optimize the use of **multiple** heterogeneous link states, L0 (ON), L2, and L3, (aka power management modes or sleeping states) already defined in ADSL/ADSL2+ standards [3]. L0 mode corresponds to the DSL being fully ON and functional, L2 mode corresponds to the case where the transmit power is a small fraction of the power at L0, whereas L3 mode corresponds to the case where no signal is transmitted on the line.
   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 this tradeoff into account, 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 lines.
2) We develop a novel **analytic framework** that computes the optimal sleeping policies for DSL lines in terms of achieving a Pareto-optimal tradeoff between energy consumption and delay performance. We take stability into

---

Manuscript received July 31, 2014; revised December 10, 2014, May 2, 2015, and September 1, 2015; accepted September 15, 2015. Date of publication September 24, 2015; date of current version December 15, 2015. The associate editor coordinating the review of this paper and approving it for publication was Prof. Andrea Tonello.

I. Kamitsos and M. Chiang are with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: kamitsos@alumni.princeton.edu; chiangm@princeton.edu).

P. Tsiaflakis is with the Department of Electrical Engineering (ESAT/SCD), KU Leuven, Leuven, Belgium (e-mail: paschalis.tsiaflakis@esat.kuleuven.be).

K. Kerpez is with ASSIA, Inc., Redwood City, CA 94065 USA (e-mail: kkerpez@assia-inc.com).

S. Ha is with the Department of Computer Science, University of Colorado, Boulder, CO USA (e-mail: sangtae.ha@colorado.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCOMM.2015.2481885

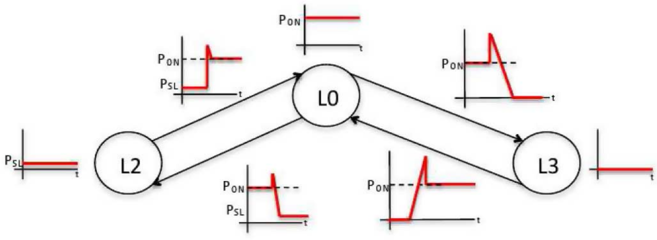

Fig. 1. Transitions between different power states (L0, L2, and L3) in DSL lines. The power consumption at L2 is non-zero but the switching cost for L2/ON transitions is smaller both in energy and time. The power consumption at L3 is zero but the switching cost for L3/ON transitions is larger both in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

account, by providing a temporal dimension solution that can significantly improve stability when using sleeping mode. We prove the hysteretic structure of the optimal sleeping policies and compute the optimal ON and OFF thresholds in terms of number of buffered jobs.

3) We extend our framework towards computing an optimal sleeping policy that **gradually** transitions a line between sub-modes of different power consumption and transmit rate and improves stability. A sub-mode has power consumption and transmit rate between that of L0 and L2/L3 operation mode. We prove the structural properties of the optimal policy, which allows the line to switch between different sub-modes, and does not require a full transition to the sleeping state (L2 or L3) before switching back on.
4) We **quantify** the benefits that arise from (a) an increased number of operating sub-modes and (b) an increased flexibility in switching between the different sub-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. In particular, we demonstrate in Section IV-A that energy savings of up to 85% can be achieved for a traffic load of 10%.
5) Our sleeping policy **results** in (a) higher energy savings compared to the Broadband Forum (BBF) TR-202 [7] policy even when the BBF's procedure is improved 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 line off more aggressively via only 6 L2 sub-modes.

A. State of the Art

The multiple power states of this research have been incorporated into the new Annex E to ITU-T G.998.4 specifying VDSL Low Power Mode (LPM) [6], particularly for stable sleeping. This paper serves as comprehensive documentation of the technical rationale and theory of operation for this new VDSL LPM. VDSL LPM saves power by transmit PSD reduction. Entry criteria into LPM states are determined by programmable settings of traffic thresholds and the time below those thresholds. A higher layer function determines the need for transition out of the L2 link states by sending primitives down to the transceivers. Exit is gradual, through multiple PSD levels, and this is configurable to allow optimization as specified in this paper. Levels of PSD reduction between steps, the maximum total PSD reduction, SNR margins, and timings between steps, are all configurable. Exit is sufficiently rapid to achieve a few Mbps in a few seconds, so the user perceives viable initial service on demand. In addition, the L2.2 link state is defined with the additional power scaling technique referred to as scheduled discontinuous operation (SDO), where the line is off for some symbol times, and while on the PSD is at a low enough level that the time-varying crosstalk does not destabilize nearby lines.

For ADSL2/ADSL2plus, the BBF guidelines in TR-202 [7] define a power saving procedure, where they recommend the use of L2 since waking up from L3 is slow and would create excessive crosstalk that could render other DSLs unstable. Specifically, the BBF TR-202 power saving guidelines are as follows:

1) It is better to use the L2 low power mode instead of L3, in order to avoid excessive crosstalk from waking up from L3.
2) The maximum aggregate transmit power reduction in L2 should not be larger than 10 dB with respect to the power consumed when the line is ON (at full power).
3) The power reduction should be gradual with steps of 1 dB apart. Together with the above guideline, this comes down to 10 steps of 1 dB, i.e., 9 intermediate states.
4) The time the line stays in each one of these intermediate states should be more than 127 seconds.

Recent research on greening DSL broadband access focuses on techniques for energy efficient dynamic spectrum management (DSM) [8], [9], [10]. They consist of the joint coordination of the DSLs' transmit spectra so as to minimize the total transmit or line-driver power consumption in the DSL access network, subject to specific QoS requirements, such as minimum data rates and fairness objectives. There has also been work on non-DSM approaches [11], [12], as well as on existing low power DSL modes [13] and rate adaptation for energy efficiency [11], [14]. This paper compliments those previous works by focusing on real-time operation of DSL low-power modes.

II. DSL SYSTEM MODEL

In Fig. 2 we display an example of a DSL access network consisting of two central offices that are connected to the Internet via two different Internet service provider networks. Multiple DSLs within the same cable bundle experience electromagnetic interactions resulting in crosstalk interference. Each central office (CO) can have a spectrum management center (SMC) to coordinate transmission over the connected DSL lines and to perform DSM. We consider a DSL system model consisting of a cable bundle of $M + N$ DSLs, in which lines $n \in \mathcal{N} = \{M + 1, \ldots, M + N\}$ are coordinated by a spectrum management center (SMC 2), whereas lines $m \in \mathcal{M} = \{1, \ldots, M\}$ are not coordinated (e.g., they transmit in the same

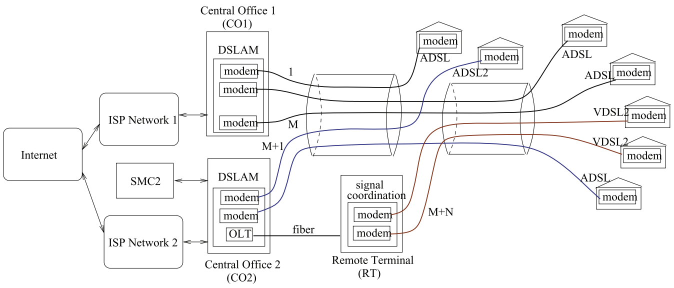

Fig. 2. Example of DSL access network with two central offices. For central office 2 there is a spectrum management center (SMC) that can coordinate the transmission of the connected DSL lines.

cable but they are managed by a different service provider). DSLs are modulated by Discrete Multi-Tone (DMT); a multicarrier transmission scheme in which individual bit streams are transmitted over independent sub-carriers, also referred to as tones. No joint signal coordination is assumed between DSLs. It is assumed that each DSL treats the crosstalk interference as additive noise. When the number of interfering DSLs is large, the crosstalk interference is well approximated by a Gaussian distribution [15]. Under this standard assumption, the resulting achievable data rate for each line n is

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

with $\text{int}_k^n = \sum_{j \neq n} h_k^{n,j} w_k^j s_k^j + \sigma_k^n$ denoting the received interference and noise of line n on tone k, $h_k^{n,n}$ the channel gain of line n on tone k, $h_k^{n,j}$ the crosstalk channel gain from line j to line n on tone k, s_k^j the transmit power of line j on tone k, σ_k^j the noise power of line j 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, see [16] Chapter 7.2. We used empirical models for the crosstalk channels that are based on 1% worst-case analysis [17].

We also include *power management mode parameters* w_k^j for each line j on tone k. We make an abstraction of the L0, L2, and L3 power management modes defined in the DSL standards. We model them by the parameters w_k^i, which can only be coordinated for the coordinated lines $\in \mathcal{N}$, i.e., fixed settings for the uncoordinated lines $w_k^j = 1, \forall k, \forall j \in \mathcal{M}$. The L0 mode of a coordinated DSL n is defined as $w_k^n = 1, \forall k$, whereas the L2 or L3 (depending which is chosen as the OFF or sleep mode) is defined as $w_k^n = \frac{P_{SL}}{P_{ON}} \forall k$, and $w_k^n = 0, \forall k$ respectively.

Note that P_{SL} corresponds to the power consumption at L2, and P_{ON} corresponds to the power consumption at L0. The power management modes of the N coordinated DSLs can be individually controlled depending on their traffic loads. The data rates and power consumptions of the coordinated DSLs depend on the chosen transmit powers and the power management modes, and are denoted as R^n and P_{ON}^n for the L0 mode of line n, respectively. The data rate of the L3 mode is zero, and the corresponding power consumption is also zero. As we focus on power management mode coordination, we assume that the transmit spectra are fixed, and they respect a minimum noise margin Γ^n for the coordinated lines when all lines are in L0.

III. OPTIMIZING OVER L2 VS L3 SLEEPING STATES

A. Model and Analysis

We assume that each line has a finite buffer B and three heterogeneous modes of operation: ON mode (L0) and two sleeping modes (L2 and L3). The power that the line 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_{ON}, whereas the transmit rate at L2 and L3 is $R_{SL} > 0$ and 0 respectively. In order to simplify the analytic derivations we will assume that $R_{SL} \approx 0$ without loss of generality, given that $R_{SL} << R_{ON}$.

In this paper, we also assume that jobs, modeling frames of packets, arrive at the line according to a Poisson process with rate λ. This essentially models data traffic such as web browsing; and while it deviates somewhat from streaming traffic the impact of streaming traffic increasing and depleting queues according to traffic load is still modeled.

The jobs have independent, identically exponentially distributed service times with mean $1/R_{ON}$. This helps analytical tractability. The memoryless property of the arrivals allows us to formulate the problem as a continuous time Markov Decision Process (MDP) problem. Note, however, that we define the traffic load ρ for each line as the ratio $\rho = \frac{\lambda}{R_{ON}}$.

The energy consumed in ON/L3 switchings is E_{ch}^3. The time required to wake up from L3 is τ_3, whereas the time required to wake up from L2 is τ_2, with $\tau_3 > \tau_2$. In addition, the cost from overhead signaling incurred during ON/L2 switchings is expressed in terms of energy as E_{ch}^2, with $E_{ch}^3 > E_{ch}^2$. In order to make the problem analytically tractable we neglect transitions between L2 and L3. The above considerations give rise to the following tradeoff: *Whenever we want to turn off the line, 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 i in the state space Ω is characterized by the triplet (W_i^2, W_i^3, Q_i), where $W_i^2, W_i^3 \in \{0, 1\}$ is an indicator that denotes whether the line was in L2 or L3 respectively during the previous time interval, and Q_i is the number of jobs in the queue. If $W_i^2 = W_i^3 = 0$ this means that the line was ON during the previous time interval. If $W_i^2 = 1$ and $W_i^3 = 0$ this means that the line was in L2 mode, where as if $W_i^2 = 0$ and $W_i^3 = 1$ this means that the line was in L3 mode. The case $W_i^2 = 1$ and $W_i^3 = 1$ is invalid.

Solving the MDP results in an optimal policy $\mathbf{p}(i) = \{p_2(i), p_3(i)\}$ for each state i, where the component $p_2(i) \in \{0, 1\}$ defines whether it is optimal or not to put the line to L2 in the next state, and the same rule applies to $p_3(i) \in \{0, 1\}$ for L3. If $p_2(i) = p_3(i) = 0$, then it is optimal to have the line ON at the next state. A case with $p_2(i) = p_3(i) = 1$ is not valid. By the Markovian structure, the policy only changes when a job arrives or departs.

Each state i evolves as $W_i^2(t+\tau) = p_2(W_i^2(t), W_i^3(t), Q_i(t))$, $W_i^3(t+\tau) = p_3\left(W_i^2(t), W_i^3(t), Q_i(t)\right)$, and $Q_i(t+\tau) = Q_i(t) + 1$, if an arrival occurs, or $Q_i(t+\tau) = Q_i(t) - 1$,

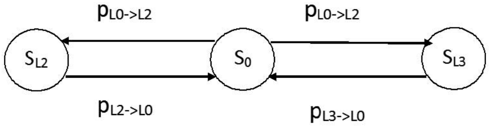

Fig. 3. Example MDP state diagram with states $S_0 = (0, 0, Q)$, $S_{L2} = (1, 0, Q + 1)$, and $S_{L3} = (0, 1, Q + 1)$. The transition probabilities between the states are depicted as $p_{L0->L2}$, $p_{L0->L3}$, $p_{L2->L0}$, $p_{L3->L0}$.

if a departure occurs. Note that t denotes the current time and $t + \tau$ refers to the time of the next event. The above equations capture the fact that the policy $\{p_2(i), p_3(i)\}$ define the mode of line switches.

From the Markovian structure of the problem, the transition probabilities from state (W_i^2, W_i^3, Q_i) are

$$Pr\Big[\big(W_i^2, W_i^3, Q_i\big) \rightarrow \big(p_2\big(W_i^2, W_i^3, Q_i\big), p_3\big(W_i^2, W_i^3, Q_i\big), Q_i+1\big)\Big] = \frac{\lambda}{\lambda + R_{ON}\left(1 - p_2\left(W_i^2, W_i^3, Q_i\right) - p_3\left(W_i^2, W_i^3, Q_i\right)\right)}, \tag{2}$$

$$Pr\Big[\big(W_i^2, W_i^3, Q_i\big) \rightarrow \big(p_2\big(W_i^2, W_i^3, Q_i\big), p_3\big(W_i^2, W_i^3, Q_i\big), Q_i-1\big)\Big] = \frac{R_{ON}\left(1 - p_2\left(W^2, W^3, Q\right) - p_3\left(W^2, W^3, Q\right)\right)}{\lambda + R_{ON}\left(1 - p_2\left(W^2, W^3, Q\right) - p_3\left(W_i^2, W_i^3, Q_i\right)\right)}. \tag{3}$$

A snapshot of the state transitions can be seen in Fig. 3. State $S_0 = (0, 0, Q)$ corresponds to the line being ON with Q jobs in the system. State $S_{L2} = (1, 0, Q + 1)$ corresponds to the line being in state L2 with $Q + 1$ jobs in the system, and state $S_{L3} = (0, 1, Q + 1)$ corresponds to the line being in state L3 with $Q + 1$ jobs in the system. The transition probabilities $p_{L0->L2}$, $p_{L0->L3}$, $p_{L2->L0}$, $p_{L3->L0}$ can be easily calculated from equations (2) and (3) by setting W_i^2, W_i^3, p_2, p_3 to appropriate values 0 or 1.

There is a cost incurred with every transition from a state to another. The so-called "cost per transition"consists of the running cost $g_1(i, p_2(i), p_3(i))$, and the switching cost $g_2(i, p_2(i), p_3(i))$ as follows:

$$\begin{aligned} g_1(i, p_2(i), p_3(i)) &= (1 - p_2(i) - p_3(i))P_{ON} + p_2(i)P_{SL} \\ &\quad + rQ_i, \qquad (4) \\ g_2(i, p_2(i), p_3(i)) &= |p_2(i) - W_i^2|E_{ch}^2 + |p_3(i) - W_i^3|E_{ch}^3 \\ &\quad + \delta|p_2(i) - W_i^2|W_i^2\lambda\tau_2 \\ &\quad + \delta|p_3(i) - W_i^3|W_i^3\lambda\tau_3. \qquad (5) \end{aligned}$$

Note that (W_i^2, W_i^3, Q_i) refers to state i.

The running cost is a weighted sum of the power consumption incurred at L0 or L2 and the delay cost incurred by jobs waiting in the queue. The congestion cost slope coefficient r models the tradeoff between power and delay cost and measures the delay sensitivity of the application in consideration, given that it captures the emphasis put on the congestion cost (i.e, number of buffered jobs). Delay sensitive applications, such as web browsing or video streaming, have relatively higher r values than delay insensitive applications, such as cloud backup. Relatively higher r values increase the emphasis put on the congestion cost and prevent the line from sleeping for too long. Then the line operates on a tradeoff point characterized by higher energy consumption and lower average delay. Relatively lower r values decrease the emphasis put on the congestion cost, and the line is allowed to sleep for longer periods. The operation point on the optimal tradeoff curve is thus characterized by decreased energy consumption and higher average delay. This parameter can also be seen as the congestion cost slope coefficient as it penalizes the number of jobs in the queue. The power cost at L0 is incurred when $p_2(i) = p_3(i) = 0$ and the power cost at L2 is incurred when $p_2(i) = 1$. Buffer overflow is not considered in this paper, since we consider only light traffic loads that model underutilized DSL lines.

The switching cost includes the energy cost incurred every time the line switches from sleeping mode (L2 or L3) to ON and vice versa. It also includes the time it takes for the line to wake up from L2 or L3, which we express in terms of the $\lambda\tau_2$ or $\lambda\tau_3$ jobs that accumulate in the buffer during the line wake up process. The energy part of the switching cost is taken equal for both turning on and off because what matters is the cost of turning on and then back off. It doesn't matter whether this is split equally between the two actions or attributed entirely to one. The scaling parameter δ models the emphasis put on the time part of the switching cost relatively to the energy part of the switching cost.

We find the policies that result in a Pareto optimal tradeoff between energy consumption and average delay, taking into account the switching cost consumed every time there is an ON → OFF or OFF → ON switching, with OFF mode corresponding to L2 or L3. Here we take into account the time it takes for the line to wake up. During the wake up process, jobs are being accumulated in the buffer, thus causing an average increase of $\lambda\tau_2$ or $\lambda\tau_3$ jobs in the queue length.

Energy is saved by reducing the number of times we expend the energy E_{ch}^2 or E_{ch}^3 required to switch on the line or turn it off to L2 or L3.

To analyze the energy-performance tradeoff, we aim at minimizing the average discounted sum of costs, as in Section 1.1.1 in [18].

$$V(i) = \min_{p_2(i), p_3(i)} \{A(p_2(i), p_3(i)) + B(p_2(i), p_3(i))\}, \tag{6}$$

where

$$A(p_2(i), p_3(i)) = \frac{g_1(i, p_2(i), p_3(i))}{\beta + v} + g_2(i, p_2(i), p_3(i)), \tag{7}$$

$$B(p_2(i), p_3(i)) = \alpha \sum_{j \in \Omega} \hat{M}_{i \to j}^{(p_2(i), p_3(i))} V(j). \tag{8}$$

where v is the uniform transition rate given by $v = \lambda + R_{ON}$, and $\hat{M}_{i \to j}$ is the uniform transition probability going from state i to state j under policy $\{p_2(i), p_3(i)\}$[1]. We apply the technique of uniformization in order to make the duration of each stage

---

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

independent of the policy chosen. The uniform transition probabilities can be calculated based on the non-uniform transition probabilities (2) and (3) as in [18].

In the objective (6) we discount future costs with a discount factor α close to 1. Discounting of future costs reflects the uncertainty of future costs, and brings the technical advantage of making the optimal cost to be finite. The optimal policy can be calculated by solving (6) using the value iteration algorithm as in Section 1.3.1 in [18]. Note that in equation (7), β is a uniformization constant given as $\alpha = v/(\beta + v)$.

The following definition is useful for the main analytic result.

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

In practice this means that if the best mode to enter with Q jobs in the queue is γ, then the system should stay there if it is already there. The proof of the following theorem is included in the Appendix of [19].

Theorem 1: The optimal policy for the objective in (6) is hysteretic.

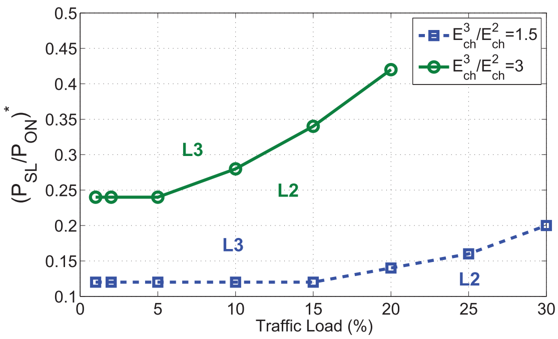

Fig. 4. Optimal ratio $(\frac{P_{SL}}{P_{ON}})^*$ vs traffic load for different values of $\frac{E^3_{ch}}{E^2_{ch}}$. For a given ratio $\frac{E^3_{ch}}{E^2_{ch}}$, increasing the traffic load results in an increased $(\frac{P_{SL}}{P_{ON}})^*$, given that the line tolerates higher values of P_{SL} in order to avoid switching from L3 too frequently. Increasing the ratio $\frac{E^3_{ch}}{E^2_{ch}}$ for a given traffic load results in the line tolerating higher values of P_{SL} without switching off to L3. Note that for each switching cost ratio, the labels L2 and L3 represent the area where it is optimal to switch off to L2 or L3 respectively.

B. *Numerical Results: L2 vs L3*

For the following numerical results we consider an ADSL2plus line with transmit power equal to 20.4 dBm and downstream data rate equal to 15.9 Mbps. These numbers are generated by using a realistic DSL simulator that was written entirely by the authors. Initially, we assume an application with $r = 10\text{J/(job*s)}$. Later in this paper, however, we sweep through a big range of r values in order to capture both delay sensitive and delay insensitive applications and trace out the Pareto optimal tradeoff between energy consumption and delay. The buffer size is taken equal to 100 and the time required to switch the line on is equal to 40s. The energy required to switch on is taken equal to the power consumption at L0 multiplied by the time required to switch the line on.We consider $\beta = 0.005$, which results in a discount factor $\alpha = 0.992$. As in [20], it is optimal to switch off the line when the buffer drains. This section examines under what conditions it is optimal to switch off to L2 versus to L3 and vice versa. The main parameters that affect this decision are the ratios $\frac{P_{SL}}{P_{ON}}$ and $\frac{E^3_{ch}}{E^2_{ch}}$. In particular, Fig. 4 shows the optimal value of the ratio $\frac{P_{SL}}{P_{ON}}$, denoted as $(\frac{P_{SL}}{P_{ON}})^*$, where the optimal sleeping state to switch off to changes from L2 to L3 or vice versa. The optimal threshold $(\frac{P_{SL}}{P_{ON}})^*$ is plotted against the traffic load for different values of the ratio $\frac{E^3_{ch}}{E^2_{ch}}$. 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 $\frac{E^3_{ch}}{E^2_{ch}}$, the increase in traffic load results in a higher $(\frac{P_{SL}}{P_{ON}})^*$. When the traffic load is higher, hence the probability of frequent switchings between ON and sleeping state (L2 or L3) increases, 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 line with a larger switching cost ratio $\frac{E^3_{ch}}{E^2_{ch}}$ 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 $\frac{E^3_{ch}}{E^2_{ch}} = 1.5$, or beyond 20% for $\frac{E^3_{ch}}{E^2_{ch}} = 3$, the optimal policy never switches the line off to L3. It instead either keeps the line always ON or switches it off to L2 when the buffer drains.

IV. STABLE SLEEPING

The introduction of L0-L2/L3 mode coordination results in time-varying crosstalk, which can lower the SNR margin of the M uncoordinated lines sufficiently to cause them to retrain, adversely impacting stability. We study the effect of varying sleep modes on the uncoordinated lines only, because stability is guaranteed for coordinated lines by design. A precise definition of stability is not straightforward and depends on many DSL implementation issues. Generally, if a line m undergoes an increase of crosstalk interference and noise that is larger than the noise margin Γ^m, it might become unstable. We define the noise impact z^m_k (for tone k and line m) as a *stability measure*, which represents the impact on the uncoordinated lines caused by the L2/L3-ON transitions of the coordinated lines:

$$z^m_k = \frac{(\text{int}^m_k)_{\text{after transition}}}{(\text{int}^m_k)_{\text{before transition}}}, m \in \mathcal{M}, k \in \{1, \ldots, K\}, \qquad (9)$$

with the denominator and numerator referring to the interference and noise on the uncoordinated lines, before and after L2/L3-ON transitions, respectively. Lower noise impact or less frequent transitions correspond to a more stable DSL network. The overall noise impact is the summation of z^m_k over all tones k and lines m.

The introduction of sleeping modes for coordinated DSLs results in an increased noise impact on the uncoordinated lines. The concern about instability induced by sleeping mode has long been a bottleneck issue. It is poorly understood and L2 mode is rarely used in current practice. This section limits the time-varying crosstalk from wake-up, which naturally enhances stability. Section V considers also gradual power reduction, as allowed by current standards.

To improve the stability behavior of the DSL network when DSLs are allowed to switch to sleeping mode, we propose a temporal dimension solution, where the wake-up of a DSL from L2/L3 to ON occurs gradually following some increasing time *function* instead of being a step function, and can be implemented by discrete, intermediate power management states, or (L2 sub-modes) with fixed time intervals equal to the wake up time, e.g., $w_k^n \in \{0, 1/8, 1/4, 1/2, 1\}, \forall k, \forall n \in \mathcal{N}$. Note that if, for example, the optimal sleeping mode is L3 and there are i L2 sub-modes, then the actual number of modes incurred during an L3 → ON switching is $i + 2$, including the L3 and ON modes. In this case L3 corresponds to an L2min sub-mode, in order to avoid confusion. Most DSLs have a 6 dB target SNR margin and they usually operate with at least 4 dB margin. If the margin falls below 0 dB they generally retrain and cause instability. So, increasing the crosstalk by a factor of two (3 dB) will rarely cause lines to retrain, but instead they can perform seamless rate adaptation (SRA) and bit-swap to adapt to the new crosstalk level.

Also, the MDP framework defined in the previous sections is applied separately for each line under consideration based on each line's particular characteristics. For the temporal dimension solution proposed here, the switching cost E_{ch} defined in the MDP framework corresponds to a full wake-up from L3 or L2 to L0. The above solution reduces the noise impact on the uncoordinated lines and thus improves the stability of these lines. The following result, which is proven in the Appendix of [19], shows how this solution improves the stability of the uncoordinated lines.

Theorem 2: Suppose that for coordinated line n there are i L2 sub-modes incurred when the line wakes up and characterized by transmit powers $\{0, \frac{1}{2^i}, \frac{1}{2^{i-1}}, \dots, \frac{1}{2}, 1\}P_{ON}$ (including the actual L2 or L3 and ON modes). Then as the number of sub-modes i increases, the noise impact on each uncoordinated line m is non-increasing.

A. Example with Stable Sleeping: Energy, Delay and Stability Performance

For the numerical evaluation of our sleeping policies through a realistic DSL simulator, we first consider an ADSL2plus scenario. We consider a situation with 12 ADSL2plus lines in downstream transmission. Lines 1–6 are uncoordinated while lines 7–12 are coordinated. The uncoordinated line lengths are 5 km, 4.8 km, 4 km, 3.8 km, 3 km, and 2.8 km, whereas the coordinated line lengths are 3 km, 2.5 km, 2 km, 1.5 Km, 1.3 km, and 1 km respectively. All the coordinated lines are located 2 km from the CO whereas all the uncoordinated lines begin at the CO. The distance from the CO has an impact on the crosstalk level: the larger this distance, the larger the crosstalk impact on the CO-connected DSL lines.

The transmit powers of the uncoordinated lines are fixed at ADSL2plus PSD mask for downstream as defined in [21]. The transmit powers of the coordinated lines are optimized based on the Distributed Spectrum Balancing (DSB) algorithm [22]. The DSM formulation used is the following:

$$\max_{\text{transmit spectra}} \quad \sum_{n \in \mathcal{N}} R^n \tag{10}$$

$$\text{subject to} \quad P^n \leq P^{n,\text{tot}}, \quad \forall n \tag{11}$$

In (10) R^n corresponds to the data rate of line n. In (11) P^n and $P^{n,\text{tot}}$ correspond to the per line n transmit power and available transmit power budget defined by DSL standards, respectively. The above conservative design guarantees stability for the coordinated lines. The design is conservative because the data rates include an SNR gap that takes a noise margin (of 6 dB) into account and in addition the system is dimensioned for the case that all lines are active and generating crosstalk. In addition, we have used the standard 1% worst-case model for the DSL channels. The optimized transmit power of each coordinated line is 0.1096 W [23]. The optimized data rates of the coordinated lines are 6.881 Mbps, 10.437 Mbps, 15.419 Mbps, 15.765 Mbps, and 16.451 Mbps respectively. The power optimization results in full allocated transmit power which is normal as the design objective for determining the transmit powers is to maximize a weighted sum of the data rates.

Given that in practice most DSLs are underutilized [24], we initially choose traffic loads below 50% in order to model such an underutilization. The 6 coordinated lines are initially assumed to operate under traffic load approximately equal to 50%, 35%, 15%, 10%, 5%, and 3%. Later, however, we will study the achievable energy savings under a wide range of traffic loads. We consider relatively delay tolerant applications and choose $r = 5\text{J/(job*s)}$. However, we later sweep through a big range of r values in order to capture both delay tolerant and sensitive applications and trace out the Pareto optimal tradeoff between energy consumption and delay performance. Assuming that the coordinated lines are characterized by $\frac{P_{SL}}{P_{ON}} = 0.2$ and $\frac{E_{ch}^3}{E_{ch}^2} = 1.5$, then the optimal sleeping state is L3. Fig. 5a presents the structure of the optimal policy of the second coordinated line for a scaling parameter of $\delta = 1000\text{J/job}$, although in Fig. 5b we calculate the optimal ON thresholds for a big range of values of δ. The ON thresholds correspond to the number of buffered jobs that a line is turned on. From Fig. 5a we observe that the optimal policy is hysteretic, as Theorem 1 predicted. The line sleeps until 9 jobs are accumulated in the buffer and then wakes up. Once awake, the line stays ON until all the jobs in the buffer are processed. Fig. 5b shows that for given value of δ, lines characterized by higher traffic loads demonstrate lower optimal ON thresholds, since they are not allowed to sleep for too long in order to process frequent incoming jobs. In addition, as δ increases, the optimal thresholds increase, given that the switching costs become higher. This means that in order to avoid excessive switching cost due to frequent switchings, all lines stay at L3 until higher queue occupancies.

Sleeping policies might cause stability issues on all uncoordinated lines, since the sudden wake up of any coordinated line produces time varying crosstalk. We will demonstrate the improvement of the stability of uncoordinated lines by exploiting the temporal solution discussed in Section IV. We consider that lines wake up according to an exponential waking up function. Specifically, when there are i L2 sub-modes the transmit powers are $\{0, \frac{1}{2^i}, \frac{1}{2^{i-1}}, \dots, \frac{1}{2}, 1\}P_{ON}$, so the power increases

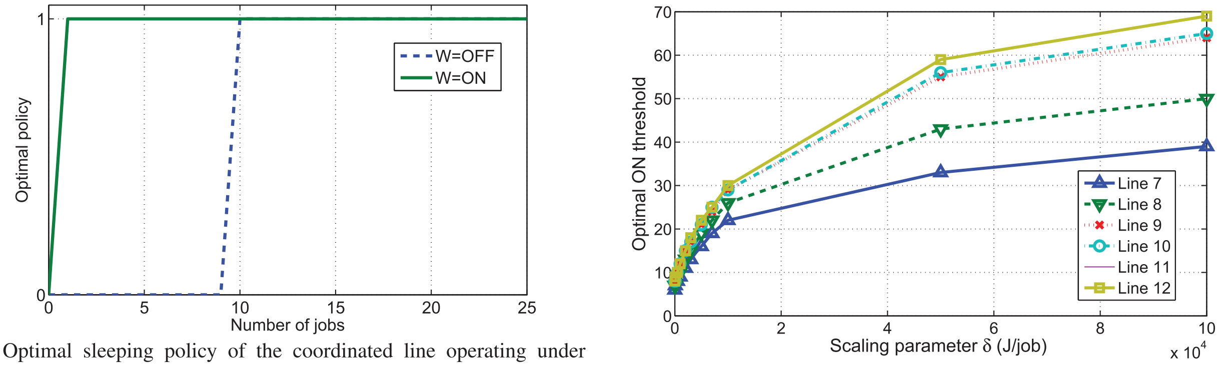

(a) Optimal sleeping policy of the coordinated line operating under traffic load 35%.

(b) Optimal ON threshold as a function of δ.

Fig. 5. Optimal policy structure

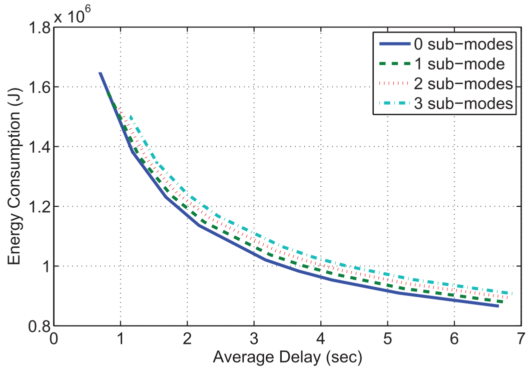

Fig. 6. Pareto optimal tradeoff between energy consumption and average delay for traffic load 10%. Making the wake up process more gradual deteriorates the energy-delay tradeoff.

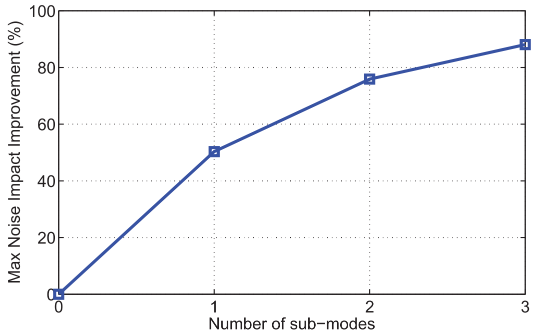

Fig. 7. Improvement of average across all uncoordinated lines of time averaged maximum noise impact with total number of L2 sub-modes incurred during L3→ON switchings. The more gradual the wake up process the bigger the improvement.

in 3 dB steps. Note that the transmit spectra of the L2 sub-modes correspond to simple power back-off versions of the full power transmit spectra obtained from (10). 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. This fixed time can be safely assumed to be lower than 1 sec.

Fig. 6 demonstrates the energy consumption - average delay Pareto optimal tradeoff for Line 10, which operates under traffic load 10%, and for different number of L2 sub-modes incurred only during L3→ON switchings. Each tradeoff curve is derived by varying the congestion cost coefficient r, with each value of this parameter corresponding to a different optimal policy. When the wake up process becomes more gradual by adding more sub-modes, the tradeoff becomes worse. This is because it takes more time for the line to fully wake up (for example, given our assumptions, one L2 sub-mode doubles the wake up time) and start transmitting data. For a given energy consumption level, the average delay increases. The increase in delay is the price we need to pay in order to reduce the noise impact on uncoordinated lines, improving the stability with respect to the direct L3→ON case. For a given average delay, the energy consumption increases with intermediate states because of the additional power consumed in each of the intermediate states. We can see that the energy consumption corresponding to arbitrarily large delay asymptotically converges to about 0.85J or 0.24 KWh. Also, delay sensitive applications, characterized by higher values of r are characterized by higher values of energy consumption and lower values of delay.

Fig. 7 shows the improvement of the averaged (across all uncoordinated lines) maximum noise impact when the number of L2 sub-modes increases. The bigger granularity of control during the wake up process can improve maximum noise impact up to 85% when there are 3 L2 sub-modes in our example. Note that the maximum noise impact per tone is important, as violating the noise impact on a tone of one line can result in errors (beyond the BER target) hence an uncoordinated line can go out of showtime. The average across all uncoordinated lines is an indicator of our policy improvement on all uncoordinated lines.

Finally, Fig. 8 shows the energy savings achieved by both the optimal sleeping policy and a baseline policy that turns a line on as soon as there is 1 job waiting in the queue to be processed. The energy savings are averaged across all lines. For all traffic loads considered, the optimal policy performs better than the baseline policy. The gap between these two policies is higher for lower traffic load, which corresponds to the realistic scenario of underutilized DSLs. Our optimal policy results in significant energy savings for lower traffic loads, i.e, around 85% for traffic load 10% and 50% for traffic load 20%. Naturally, the energy savings decrease as the traffic load becomes higher, which reduces the lines' sleeping times.

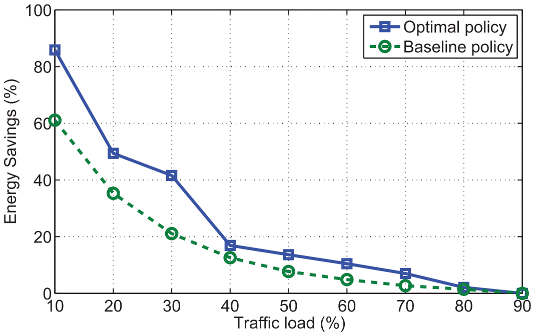

Fig. 8. Energy savings over traffic load for optimal and baseline sleeping policies. The optimal policy performs better especially for lower traffic loads.

V. COMPARING BBF TR-202 GUIDELINES AND OUR NEW DESIGN

The BBF TR-202 guidelines [7] described in I-A have MDP cost per transition that takes the form

$$\begin{aligned}\hat{g}(i, p(i)) = &\frac{1}{\beta + v}(p(i)P_{ON} + (1 - p(i))P_{SL} + rQ_i)\\ &+ |W_i - p(i)|E_{ch}^2 + |p(i) - W_i|W_i \sum_{l=1}^{9} P_l \tau_{\text{off},l}\\ &- \delta|p(i) - W_i|W_i \sum_{l=1}^{9} R_{\text{off},l}\tau_{\text{off},l} \qquad (12)\end{aligned}$$

In the above formula, P_k and $R_{\text{off},k}$ denote the power consumption and data rate in intermediate state k introduced by the above guidelines respectively. The time spent in each of the intermediate states is $\tau_{\text{off}} = 127s$.

Note that (12) models the case in which the line gradually goes to the lowest L2 mode each time it decides to save power, which corresponds to a worst case delay and minimum power configuration, whereas in practice the line can stop at an intermediate level from which it can jump back to full power mode. Also (12) increases the switching cost, when the line turns off, by the amount of energy spent in each of the intermediate states. At the same time the delay component of the switching cost is decreased because jobs can still be processed while the line is in these intermediate states.

In this section, we leverage the analytic framework introduced in Section III-A and propose a scheme that takes into account the gradual power-down and power-up of the line, providing for an implementation using intermediate states or operating sub-modes. Here we introduce operating sub-modes for both waking up or going to sleep. We seek to find the optimal policy that controls such transitions, taking into account the fact that once a line is in the process of waking up, or going to sleep, it does not need to transition through all the sub-modes. Also, in contrast to the previous approach in [4], a line does not need to switch off all the way to the sleeping state before switching back on again. This provides more flexibility.

A. Analytic Framework

As in Section III-A, 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 with $c + 1$ L2 sub-modes $(0, 1, 2, \ldots, c)$, with the mode c referring to full power state L0. The sub-mode 0 refers to the sleeping state (L2 or L3, depending on which sleeping state is optimal) and is alternatively called L2min or L3min sub-mode. The decision whether the sleeping state is chosen to be L2 or L3 depends on the modem characteristics and is based on the analysis in Section III-A. We also consider that each line 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 $i = (W_i, Q_i)$ if the line was in sub-mode W_i in the *previous* time interval and *currently* there are Q_i jobs in the system.

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

The solution to the MDP is a *policy* denoted by p, indicating the optimal sub-mode (the action) for the line to be in each MDP state. The transition probabilities at state (W_i, Q_i) with policy $p(W_i, Q_i)$ and $0 < Q_i < B$ now become:

$$Pr[(W_i, Q_i) \to (p(W_i, Q_i), Q_i + 1)] = \frac{\lambda}{\lambda + R_m p(W_i, Q_i)}, \qquad (13)$$

$$Pr[(W_i, Q_i) \to (p(W_i, Q_i), Q_i - 1)] = \frac{R_m p(W_i, Q_i)}{\lambda + R_m p(W_i, Q_i)}. \qquad (14)$$

The running and switching costs, g_1 and g_2 respectively, now become:

$$g_1(i, p(i)) = (p(i)P_{\text{m}} + \mathbf{1}_{p(i)=0}P_{OFF} + rQ_i), \qquad (15)$$

$$g_2(i, p(i)) = |W_i - p(i)|e_s + \delta|W_i - p(i)|\lambda\tau_s. \qquad (16)$$

Here e_s and τ_s represent the energy and time required for switching between two successive sub-modes. Note that $P_{OFF} = P_{SL}$, if the optimal sleeping state is L2min, and $P_{OFF} = 0$, if the optimal sleeping state is L3min.

The parameter δ, expressed in J/s, models the emphasis put on the time part of the switching cost relatively to the energy part.

Our new sleeping mode operation design can be summarized as follows.

1) The optimal sleeping state can be either L2 or L3 based on the line's characteristics. The optimization decision is based on the analysis of Section III-A.
2) Introduce multiple L2 sub-modes with power consumption and transmit rate $W_i P_m$ and $W_i R_m$ respectively, where W_i is the W_i-th sub-mode above L2min or L3min, $P_m = P_{ON}/c$ and $R_m \approx R_{ON}/c$.

3) The line can visit any sub-mode in an increasing or decreasing order while turning on or off respectively, and does not necessarily need to go through all sub-modes while waking up or going to sleep. The sequence of sub-modes to visit depends on the optimal policy.
4) There is a cost in terms of energy and time for switching between different sub-modes. This costs are denoted as e_s, and τ_s for successive sub-mode switching respectively.
5) Switching between different sub-modes takes place on an event of an arrival or departure and only if it is dictated by the optimal policy.

In addition, and contrary to the BBF TR-202 guidelines, our scheme involves the use of a buffer where jobs are being accumulated, so the number of jobs that need to be processed is readily known. From a practical implementation perspective, a buffer can correspond to some amount of SDRAM memory that needs to be carefully dimensioned [25] [26].

Theorem 3: The policy that optimally decides the transitions between different sub-modes is monotone hysteretic.

The proof of this theorem is provided in the Appendix of [19]. A formal definition of the monotonicity of a hysteretic polity can be found in [27]. In our problem the monotone hysteretic policy is characterized by $2c$ thresholds, c for switching sub-modes while the line is turning on, and c for switching sub-modes while the line is turning off.

B. Simulation Results: New Scheme

1) Optimal Policy Structure: We illustrate the gradual transition design in an ADSL2plus 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 (RT) 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 20.4 dBm. The coordinated lines are initially under traffic loads 50%, 35%, 15%, 10%, 5%, and 3% respectively.

We pick those traffic loads to model underutilized DSLs (with the exception of 50% traffic load). As mentioned earlier, such light traffic loads allow to not consider buffer overflow in this paper. Initially, we consider $c = 5$ (6 sub-modes), however we later study the system performance under a large range of sub-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 $\frac{E_{ch}^3}{E_{ch}^2} = \frac{3}{2}$. From Fig. 4 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 line to L3. Note that in our framework L2 corresponds to sub-mode L2min. We also consider a delay insensitive application with $r = 10$J/job. Finally, we assume that the time required to transition between successive sub-modes is 8s, resulting in a total wake up time of 40s, and $\delta = 1$J/s, thus putting equivalent emphasis on the time and energy part of the switching cost. The mean interarrival time is taken equal to 10s, thus making the arrival of a job during mode-transitioning rather not probable.

Fig. 9 demonstrates the optimal policy for line 4. We can see that the optimal ON thresholds for switching to a higher submode are equal to 2, 6, 9, 13, and 16 jobs whereas the optimal OFF thresholds for switching to a lower submode are equivalent to 4, 3, 2, 1, 0 jobs. Note that the traffic load of 35%, and the delay-insensitive application does not require the line to be fully ON until there are 16 jobs in the buffer.

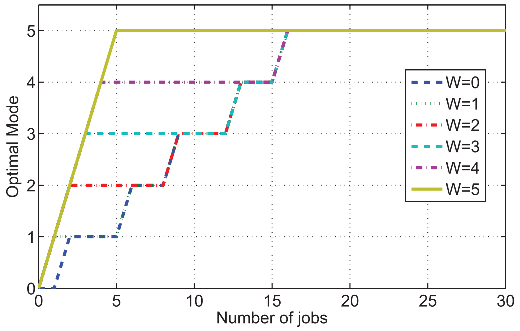

Fig. 9. Optimal policy for $c = 5$ and traffic load of 35%. The optimal sleeping state is L2.

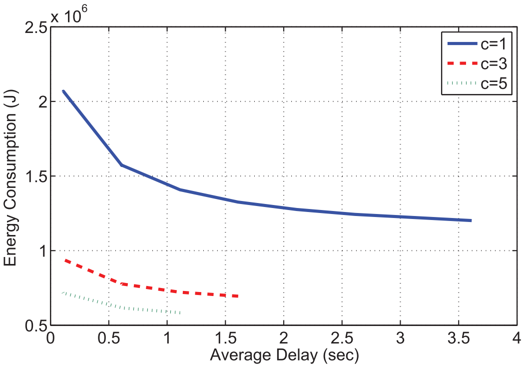

Fig. 10. Pareto optimal tradeoff between energy consumption and average delay for the coordinated line 4 under our optimal sleeping policy (c=5). Our approach results in a better energy-delay tradeoff when the number of L2 sub-modes available increases.

In this particular case the optimal policy suggests that the modem should visit all the intermediate sub-modes in order to fully turn on. Switching immediately to full power would be suboptimal in terms of balancing energy consumption and delay performance. However, if the designer of the line still wishes a suboptimal operation where the line switches on directly to full power from any L2 sub-mode, an option would be to use 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 L2 sub-modes. In particular, Fig. 10 demonstrates the Pareto optimal tradeoff between energy consumption and average delay for the coordinated line 4 (traffic load equal to 10%) in two cases; when there are 2 sub-modes available ($c = 1$) and when there are 6 sub-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 is reduced by around 50-65% and average delay is reduced by 5-65% for all delay sensitivities considered.

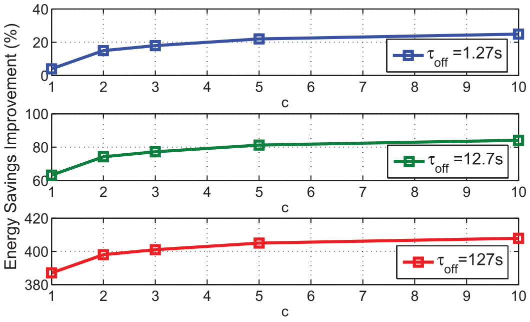

Fig. 11. Improvement of average energy savings across all coordinated lines achieved by our optimal policy for different numbers of L2 sub-modes in comparison with the BBF policy under different values of τ_{off}.

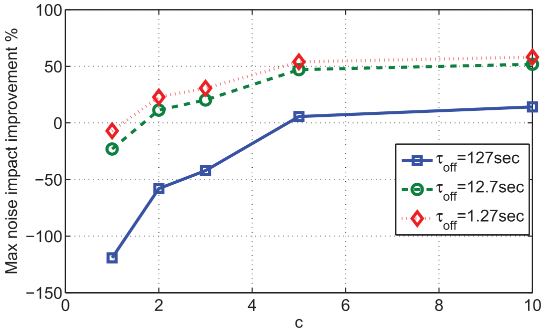

Fig. 12. Improvement on maximum noise impact achieved by our sleeping policy compared to the BBF guidelines approach for different values of τ_{off}. As τ_{off} decreases, our approach results in better stability for a large range of L2 sub-modes.

In addition, Fig. 11 shows the improvement of energy savings achieved by our scheme for different number of L2 sub-modes in comparison with those achieved by the BBF policy as described earlier and modeled in equation (12). We evaluate a modified version of the guidelines in the BBF TR-202 that allows different values of τ_{off} to be spent in each intermediate state.

Due to higher flexibility of control, as the number of L2 sub-modes increases, the improvement of energy savings achieved by our optimal scheme increases as well. BBF TR-202 guidelines perform poorly when the time spent in each intermediate state before turning off is high ($\tau_{off} = 127$s). In particular, our scheme performs almost 4 times better depending on the number of L2 sub-modes. When τ_{off} decreases, the performance improvement naturally falls and ranges from 63% to 84% for $\tau_{off} = 12.7$s, and from 4% to almost 25% for $\tau_{off} = 1.27$s, depending on the number of L2 sub-modes chosen.

3) Stability: Fig. 12 illustrates the improvement of the maximum noise impact (maximum z_k^m for all tones k and lines m) achieved by our optimal sleeping policy for different numbers of L2 sub-modes in comparison with the maximum noise impact achieved by the modified BBF TR-202 guidelines' sleeping method having 9 L2 sub-modes.

When the time spent in each of the intermediate states defined by the BBF TR-202 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 system stability. 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 TR-202 guidelines reduces the sudden increase of electromagnetic coupling in the bundle relatively well. The original BBF TR-202 approach with $\tau_{off} = 127$s however, performs much worse in terms of energy savings compared to our optimal scheme. Note that all comparisons have been made under the consideration that the BBF TR-202 guidelines use all 9 intermediate states for good stability performance.

VI. SUMMARY

The use of DSL sleeping modes offers great potential for decreasing energy consumption of the broadband access infrastructure. In this paper, we develop optimal and stable sleeping mode policies. We provide an analytic framework to optimally decide between L2 and L3 sleeping modes, based on the ratio of the respective switching costs and the power consumption at L2. It consists of a Markov Decision Process (MDP) formulation with easily tunable parameters that model DSL sleeping modes. Our MDP-based optimal policies result in a Pareto-optimal tradeoff between energy consumption and delay performance. We demonstrate that compared to a baseline policy, which turns a DSL on as soon as there is an incoming job to be transmitted and turns it back off as soon as the job has been transmitted, our optimal policy results in higher energy savings for all traffic loads. The performance improvement becomes higher for lower traffic loads which correspond to the realistic scenario of underutilized DSLs.

The impact of the optimal policy on stability is furthermore studied by considering the impact of time-varying crosstalk that could destabilize uncoordinated DSLs. A stability improving mechanism is then proposed so as to obtain stable sleep mode operation. It consists of smoothing the wake-up process of DSLs in time by introducing intermediate power states, or L2 sub-modes. The benefits of this mechanism are evaluated by numerical simulation results, demonstrating a three-way tradeoff between energy consumption, delay performance, and stability.

In addition, we prove the monotone hysteretic structure of the optimal policy that transitions a DSL between different submodes in order to improve the Pareto optimal tradeoff, without requiring the line to completely switch off before turning back on again. We compare the performance of our approach with the current BBF TR-202 sleeping guidelines, and 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 TR-202 policy is modified to stay in each of the intermediate states a much shorter amount of time. In addition, our numerical results show improved stability performance of our policy compared to the BBF TR-202 guidelines, even when the BBF TR-202 policy operates with 127s in each of its 9 intermediate states and our policy considers only 6 L2 sub-modes. The stability improvement increases to more than 50% when our policy considers 11 L2 sub-modes ($c = 10$), which is equivalent to 9 intermediate states in BBF TR-202 guidelines, and BBF TR-202 policy operates with 1.27s in each intermediate state.

To implement this in practice, ADSL2/ADSL2plus low-power link states operation can be controlled to enable the multiple sub-states and gradual entry as described here. For VDSL2, at this time of writing, the ITU-T has recently completed a new Annex specifying VDSL2 Low Power Mode (LPMode) [6]. The multiple power states and gradual exit advanced in this paper have been incorporated into this Annex, which can be directly implemented for enhancing energy efficiency and stability of VDSL2. This new VDSL2 LPMode allows gradual entry and exit through multiple low power states. This gives DSL operators the ability to implement the theory advanced in this paper to optimize power usage with minimal delay and while maintaining stability.

REFERENCES

[1] O. Johnson. (Apr. 2013). [Online]. Available: http://point-topic.com/press-and-events/2013/the-dsl-king-is-dead-long-live-the-vdsl-king

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

[3] G. Ginis, "Low-power modes for ADSL2 and ADSL2+," Broadband Communications Group, Texas Instruments, White paper SPAA021, Jan. 2005.

[4] I. Kamitsos, P. Tsiaflakis, S. Ha, and M. Chiang, "Stable sleeping in DSL broadband access: Feasibility and tradeoffs," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Houston, TX, USA, Dec. 2011, pp. 1–6.

[5] I. Kamitsos, P. Tsiaflakis, K. Kerpez, S. Ha, and M. Chiang, "Energy efficient DSL via heterogeneous sleeping states: Optimization structures and operation guidelines," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, 2012, pp. 3128–3134.

[6] Amendment. (Jan. 2015). Improved impulse protection for DSL transceivers 4 to Recommendation ITU-T G.998.4, 2010.

[7] "ADSL2/ADSL2plus low-power mode guidelines," Broadband Forum TR-202, Feb. 2010.

[8] J. M. Cioffi *et al.*, "Greener copper with dynamic spectrum management," in *Proc. 3rd Int. Conf. AccessNets*, Las Vegas, NV, USA, Oct. 2008, pp. 1–5.

[9] P. Tsiaflakis, Y. Yi, M. Chiang, and M. Moonen, "Fair greening of broadband access: Spectrum management for energy-efficient DSL networks," *EURASIP J. Wireless Commun. Netw.*, vol. 2011, p. 140, 2011.

[10] K. Hooghe and M. Guenach, "Toward green copper broadband access networks," *IEEE Commun. Mag.*, vol. 49, no. 8, pp. 87–93, Aug. 2011.

[11] K. Hooghe and M. Guenach, "Towards energy-efficient packet processing in access nodes," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Houston, TX, USA, Dec. 2011, pp. 1–6.

[12] K. Hooghe and M. Guenach, "Impact of FTTN architecture on access node energy efficiency," in *Proc. 17th IEEE Symp. Commun. Veh. Technol. Benelux (SCVT)*, 2010, pp. 1–5.

[13] M. Guenach *et al.*, "Power-efficient copper access," *Bell Labs Tech. J.*, vol. 15, no. 2, 2014, pp. 1–5.

[14] M. Guenach, C. Nuzman, J. Maes, and M. Peeters, "Trading off rate and power consumption in DSL systems," *Proc. IEEE GLOBECOM Workshop Green Commun. (GreenCom2)*, Honolulu, HI, USA, Nov./Dec., 2009, pp. 117–129.

[15] K. J. Kerpez, "Near-end crosstalk is almost Gaussian," *IEEE Trans. Commun.*, vol. 41, no. 5, pp. 670–672, May 1993.

[16] T. Starr, J. M. Cioffi, and P. Silverman, *Understanding Digital Subscriber Line Technology*. Upper Saddle River, NJ, USA: Prentice-Hall, 1999.

[17] ETSI, "Very high speed digital subscriber line (VDSL); Functional requirements," ETSI Standard TS 101 270-1, Rev.V.1.3.1, 2003.

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

[19] I. Kamitsos, P. Tsiaflakis, K. Kerpez, S. Ha, and M. Chiang. (2015). "Stable sleep mode optimization for energy efficient DSL," Tech. Rep. [Online]. Available: http://www.princeton.edu/~chiangm/DSL_stable_sleep.pdf

[20] I. Kamitsos, L. Andrew, H. Kim, and M. Chiang, "Optimal sleeping patterns for serving delay tolerant jobs," in *Proc. e-Energy*, Passau, Germany, Apr. 2010, pp. 31–40.

[21] Annex A. (Jan. 2009). Asymmetric digital subscriber line (ADSL) transceivers - extended bandwidth ADSL2 (ASDL2plus) of ITU-T G.992.5.

[22] P. Tsiaflakis, M. Diehl, and M. Moonen, "Distributed spectrum management algorithms for multiuser DSL networks," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 4825–4843, Oct. 2008.

[23] ANSI, "Spectrum management for loop transmission systems," T1E1.4/2003-002R1, ANSI draft, no. 2, 2003.

[24] E. Goma *et al.*, "Insomnia in the access or how to curb access network related energy consumption," *Proc. ACM SIGCOMM*, Toronto, ON, Canada, 2011, pp. 338–349.

[25] Intel. "Digital subscriber line access multiplexer (DSLAM) example design," Intel application note, May 2002.

[26] C. Sauer, M. Gries, and S. Sonntag, "Modular reference implementation of an IP-DSLAM," in *Proc. 10th IEEE Symp. Comput. Commun. (ISCC)*, Jun. 2005, pp. 191–198.

[27] S. K. Hipp and U. D. Holzbaur, "Decision processes with monotone hysteretic policies," *Oper. Res.* vol. 36, no. 4, pp. 585–588, Jul./Aug. 1988.

[28] F. V. Lu and R. F. Serfozo, "M/M/1 queueing decision processes with monotone hysteretic optimal policies," *Oper. Res.*, vol. 32, no. 5, pp. 1116–1132, Sep./Oct. 1984.

Ioannis Kamitsos received the Diploma degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 2006, and the M.A. and Ph.D. degrees from Princeton University, Princeton, NJ, USA, in 2009 and 2012, respectively. In the summer of 2010, he interned at the Standards and RF Laboratory, Samsung Telecommunications America, Richardson, TX, USA. His research interests include optimization theory, optimal control and their applications on energy efficient computing, communications, networks, and Smart Grid.

Paschalis Tsiaflakis (S'06–M'09) received the Master's degree in electrical engineering and the Ph.D. degree in engineering sciences from the KU Leuven University, Leuven, Belgium. From 2011 to 2013, he was a Postdoctoral Researcher, during which he was a Visiting Researcher at Princeton University, Princeton, NJ, USA; UCLA, Los Angeles, CA, USA; Tsinghua University, Beijing, China; and UCLouvain, Louvain-la-Neuve, Belgium. Currently, he is a Researcher with Bell Labs Alcatel-Lucent, where he focuses on innovation in next generation ultrabroadband access. He has authored over 50 journal and conference papers. His research interests include optimization and communication theory for wireline and wireless communication systems. He was the recipient of the Belgian Young ICT Personality Award in 2010, two IEEE ICC Best Paper Awards in 2013, and was a top-12 finalist for the European ERCIM Cor Baayen Award 2010.

Kenneth J. Kerpez (M'89–SM'01–F'04) received the Ph.D. degree from Cornell University, Ithaca, NY, USA, in 1989. He worked with Bellcore and Telcordia for 20 years and has been working with ASSIA for 5 years. His research interests include communications systems and networks of all sorts, including DSL, fiber access, home networks, wireless systems, broadband service assurance, IPTV, IP QoS, triple-play services, and virtualization.

Sangtae Ha (S'07–M'09–SM'12) received the Ph.D. degree in computer science from North Carolina State University, Raleigh, NC, USA. He is an Assistant Professor with the Department of Computer Science, University of Colorado at Boulder, Boulder, CO, USA. He cofounded the Princeton EDGE Lab as its first Associate Director in 2009 and led its research team as an Associate Research Scholar with Princeton University, Princeton, NJ, USA, from 2010 to 2013. He is a Cofounder and the founding CTO/VP Engineer of DataMi, a startup company on mobile networks, and is a Technical Consultant to a few startups. His research interests include building and deploying practical network systems. He serves as an Associate Editor for IEEE INTERNET OF THINGS JOURNAL. He was a General Cochair of the workshop on Fog Networking for 5G and IoT in IEEE SECON 2015 and the Smart Data Pricing workshop in IEEE INFOCOM 2013. He was the recipient of the INFORMS ISS Design Science Award in 2014.

Mung Chiang (S'00–M'03–SM'08–F'12) is the Arthur LeGrand Doty Professor of Electrical Engineering at Princeton University, Princeton, NJ, USA. His research on networking received the 2013 Alan T. Waterman Award. His textbook, *Networks: Friends, Money and Bytes,* and online course reached 250,000 students since 2012. He founded the Princeton EDGE Lab in 2009, which bridges the theory-practice gap in edge networking research by spanning from proofs to prototypes, and co-founded a few startups in mobile, IoT, and big data. Professor Chiang is the Director of Keller Center for Innovations in Engineering Education at Princeton University and the inaugural Chairman of Princeton Entrepreneurship Council.