Performance Evaluation of Power Management for Best Effort Applications in IEEE 802.16 Networks

G. Anastasi[1], A. Bacioccola[2], C. Cicconetti[1], L. Lenzini[1], E. Mingozzi[1], Carlo Vallati[1]
[1]Dipartimento di Ingegneria dell'Informazione, University of Pisa, Italy
[2]Nokia Devices, Helsinki, Finland.
{g.anastasi, c.cicconetti, l.lenzini, e.mingozzi, c.vallati}@iet.unipi.it, a.bacioccola@nokia.com

Abstract—**Reduced energy consumption is a crucial aspect of mobile Broadband Wireless Access (BWA) networks, which are expected to be populated by battery-operated devices, like mobile phones and palmtops. For this reason the IEEE 802.16e, which is one of the front-runner competitors in this field, specifies a set of power saving mechanisms to be employed by the Base Station (BS) and Mobile Stations (MSs) to reduce the amount of time the latter spend with the wireless interface on. These mechanisms are classified into three class types, which are designed for different types of applications. In this paper we focus on class type I, which fits the typical requirements of best-effort traffic. With class type I, an MS with power saving enabled alternates between sleep and listening periods. The duration of the sleep periods increases by a factor 2 each time a listening period ends, up to a maximum sleep window size. Since the standard does not provide guidelines for setting the above parameters, which are negotiated between the BS and MSs when setting up a power saving class, we evaluate via simulation their impact on the performance, in terms of both application-specific metric, i.e. delay or throughput depending on the type of traffic, and the amount of energy saved.**

I. INTRODUCTION

Mobile Broadband Wireless Access (BWA) is one of the hottest areas for research and development of new technologies in the context of wireless networks. It enables users to have the same broadband experience in a ubiquitous manner. One of the strongest competitors among mobile BWA technologies is IEEE 802.16e [1], which has been ratified in 2005 and was designed to support a wide variety of QoS-enabled and/or capacity demanding applications, thus extending broadband access to mobile users. However, unlike fixed broadband access technologies, support for an efficient energy resource management on the terminal side is also important because mobile devices (e.g., notebooks, PDAs, cellular phones) typically have limited energy. For this reason, a *sleep mode* is included in the IEEE 802.16 standard. When a terminal enters the sleep mode, it switches off the communications capabilities of its wireless network adapter. Efficient support of sleep mode has been recognized as one of the most important factors in the definition of the IEEE 802.16m standard, which is currently being defined to further improve the features of IEEE 802.16e.

In this paper, we carry out a performance analysis of the power saving mechanism in IEEE 802.16e through simulation. We consider interactive and non-interactive applications separately because they have very different requirements from the point of view of power saving. In fact, interactive applications continuously require input from a user, even though at different time scales depending on the specific application. Examples of interactive applications include instant messaging, online gaming and web browsing. On the other hand, non-interactive applications, like file transfer, do not have specific time constraints and are expected to run in background with respect to the normal user activities.

The paper is organized as follows. Section II provides a brief introduction to the IEEE 802.16 standard, which also reviews the related work. Section III includes the performance analysis. The simulation assumptions and metrics definition can be found in the same section. Conclusions are drawn in Section IV.

II. POWER SAVING MECHANISMS IN IEEE 802.16e

In this section, we provide a brief introduction to the IEEE 802.16 standard. We only focus on those aspects which are relevant to our analysis. A survey on the IEEE 802.16 MAC protocol can be found in [2], while we refer to the standard document [1] for a comprehensive description of the power saving mechanisms.

The IEEE 802.16 standard defines a Point to Multi-Point (PMP) network where the BS broadcasts to all the MSs in the *downlink* direction and it centrally coordinates the transmission from the MSs in the *uplink* direction. The Medium Access Control (MAC) protocol is connection-oriented. Each connection is uniquely identified by a connection identifier (CID). Connections are uni-directional, i.e. either uplink or downlink, and MSs can establish multiple connections with the BS. The IEEE 802.16 MAC protocol is frame-based, i.e. time is partitioned into frames of fixed duration.

The IEEE 802.16 standard supports power management by means of the *sleep mode*. When an MS enters the sleep mode it is said to be in a *sleep window*, and it cannot communicate with the BS: no uplink or downlink traffic can be received from/addressed to the MS. An MS that is not in a sleep window is said to be in a *listening window*, because it receives broadcast transmissions from the BS. How to enter/leave the sleep mode is negotiated through a set of messages between

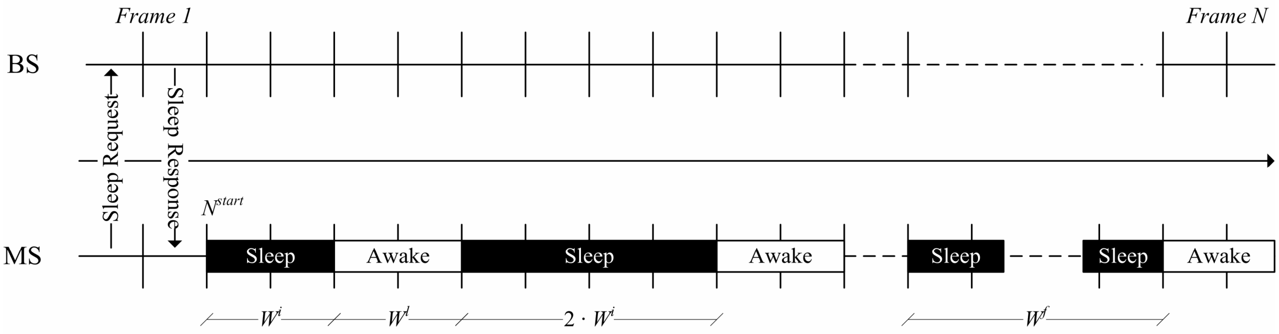

Fig. 1. Power saving class of type I.

TABLE 1
RELEVANT POWER SAVING PARAMETERS (IN FRAMES)

Name	Symbol
Initial-sleep window size	W^i
Listening window size	W^l
Final-sleep window size	W^f
Start frame	N^{start}

the MS and the BS. In order for the BS not to send data to an MS during its sleep windows, the BS is required to keep one *power saving sleep class* (or *class* only), which is uniquely identified by a sleep ID (SLPID). One or more connections with common power saving demand properties can be associated to a single SLPID, while a CID can only be mapped to one SLPID. It follows that an MS is only allowed to switch off its radio transceivers when all its connections belong to classes in their sleep windows. For ease of readability we assume without loss of generality that each MS only has a single class to which all its connections are mapped. Therefore, the MS switches off its transceivers if and only if its class is active and in a sleep window.

The standard defines three different types of classes: type I, type II, and type III, respectively. These three class types differ on the set of parameters used to identify a class and on the procedures for alternating between sleep and listening windows. Specifically, class type I is targeted to support power saving for best effort and non-real-time variable bit-rate applications. Class type II is mainly targeted to support power saving for real-time applications, and finally class type III is intended for supporting power saving for multicast and management connections. In this work we focus on class type I only, which is described in the following. Therefore, in the remainder of this work we always refer to class type I.

When a class becomes active, it alternates between sleep and listening periods. If the MS leaves the sleep mode, then the class is de-activated, and the MS remains in a listening state. Listening windows have fixed duration, equal to W^l frames, called *listening window size*, while the duration of sleep windows increases from a minimum value of W^i frames, called *initial sleep window size*, to a maximum value of W^f frames, called *final sleep window size*. Specifically, the sleep window size doubles after each listening window until the maximum is reached. The first sleep window begins as defined by the parameter N^{start}, in frames. The relevant power saving parameters are reported in Table 1 and illustrated in Fig. 1.

A class can be initiated by either an MS or the BS, according to the following rules. In the MS initiated procedure, the MS sends a *sleep request* (MOB_SLP-REQ) message to the BS containing the set of the preferred power saving parameters. The BS can accept, refuse or change those parameters in a *sleep response* (MOB_SLP-RSP) message. In the BS initiated procedure an *unsolicited sleep response* message is sent by the BS to the MS. In this case the parameters cannot be re-negotiated by the MS. A class can be de-activated by either the BS or an MS. On the one hand, the MS can deactivate the class when uplink traffic is detected. In this case, de-activation is notified to the BS through a special MAC header which is sent to the BS in a contention manner, and also includes the amount of data waiting for transmission for the connection that became backlogged. On the other hand, the BS can de-activate a class only during a listening window of the MS. This is done by sending a positive *traffic indicator* (MOB-TRF-IND) message to the MS during a listening window. Traffic indicator messages are broadcasted by the BS and contain a list of positive or negative indication to the MSs. If the indication is positive, then there are data waiting for transmission to the MS, which de-activates the class. Otherwise, the MS continues to alternate between sleep and listening windows.

A. Related Work

In the last years, research in power saving techniques has focused mainly in technologies for WLANs, especially IEEE 802.11, due to its large diffusion and popularity. The performance of the IEEE 802.11 power saving mode has been investigated in detail, and several enhancements and solutions have been proposed [3, 4, 5]. However, the existing solutions cannot be straightforwardly applied to IEEE 802.16, due to the inherent properties of the power saving mechanism defined in this context.

On the other hand, the literature so far lacks substantial work with specific reference to IEEE 802.16. Most existing studies, e.g. [6, 7, 8], propose mathematical models of power saving in IEEE 802.16e, which by necessity introduces several assumptions about the system and traffic characterization. For instance, Poisson is typically assumed as the traffic arrival process, which however is known to capture poorly the characterization of Internet traffic (e.g. [9]). Instead, in [10] the

TABLE 2
POWER SAVING PARAMETER VALUES

Name	Symbol	Range of values (frames)
Initial-sleep window size	W^i	[1, 50]
Listening window size	W^l	[1, 50]
Final-sleep window size	W^f	{128, 256, 1024}
Idle time	T^{idle}	[1, 50]

authors consider Constant Bit Rate (CBR) and File Transfer Protocol (FTP) via simulation analysis. Their results show that with Transmission Control Protocol (TCP) traffic the performance of power saving mechanisms directly depends on the Round Trip Time (RTT) experienced: the higher the RTT is, the higher the transfer delay becomes, but the lower the amount of energy consumed by MSs becomes.

III. PERFORMANCE ANALYSIS

In this section, we evaluate the system performance of IEEE 802.16, in terms of both application metrics, such as throughput and delay, and energy-related metrics, like the activity cycle. Results for interactive and non-interactive traffic are presented separately.

The simulations were carried out by means of a prototypical simulator of the IEEE 802.16 MAC protocol. The simulator is event-driven and was developed as a module for the popular simulator ns-2 [11]. The simulation experiments have been carried out using the method of independent replications [12], using the framework for statistics gathering in ns-2 described in [13]. The MAC layer of the BS and MSs are implemented, including all procedures and functions for uplink/downlink data transmission and uplink bandwidth request/grant. Both the MS and BS power saving state machines are implemented and the simulator provides the support for dynamic power saving class activation/de-activation.

In our analysis we assume that a traffic indicator message is always sent by the BS to any MS with at least one class active during its listening window. Additionally, we introduce a parameter, called *idle time* (T^{idle}), which is used to detect inactivity by both the BS and MSs as follows. Whenever a MAC Service Data Unit (SDU) belonging to a connection of a SLPID is received at the MAC layer, a timer is started equal to T^{idle} frames for that SLPID. If another instance of the same timer is running, then it is reset to the initial value T^{idle}. When the timer expires it is assumed that the connections of that SLPID will remain idle for some time, hence a power saving class is activated. The performance is evaluated by varying all the power saving parameters in the range reported in Table 2.

In all our simulation experiments system parameters are configured as reported in Table 3, which is based on the system profile specified by the WiMAX Forum in [14]. Specifically, the network is simulated under ideal channel conditions: all PHY Payload Data Units (PDUs) are delivered without errors from the BS to the MSs and vice versa. This allows us to derive an upper bound of the performance obtained for each simulated traffic type, and in turn assess the effectiveness of the power saving mechanism specified by the IEEE 802.16e standard, with no bias due to the specific radio propagation model employed. Therefore, error recovery mechanisms at the MAC and PHY layers are disabled. Furthermore, all the results have been obtained with MSs employing the same modulation and coding scheme (MCS), i.e. QSPK-3/4, during the whole simulation duration, while the control messages were transmitted with QPSK-1/2 repetition code 4. The former conveys 9 bytes per slot, while the latter is much more robust since any slot only carries 1.5 bytes.

TABLE 3
PHYSICAL / MAC PARAMETER VALUES

Name	Value
Duplexing	TDD
DL Sub-channelization scheme	PUSC
UL Sub-channelization scheme	PUSC
Physical channel model	Ideal
OFDM symbol duration	106.382 µs
Number of sub-channels	30
Number of OFDM symbols per frame	47
Downlink sub-frame / uplink sub-frame	2.91
Frame duration	5 ms
Uplink allocation start time	8.723 ms
MAC data MCS	QPSK-3/4
MAC control MCS	QPSK-1/2 rep.code 4

To asses the performance of interactive data applications we use the following metrics: *delay*, *throughput*, and *duty cycle*. The *delay* is defined as the time interval between when the first bit of a MAC SDU is received at the BS/MS MAC layer and when the last bit of the same MAC SDU is received at the MS/BS MAC layer. The *throughput* is the number of bits correctly received by an MS in the unit of time. Finally, the *duty cycle* is defined as the ratio between the total time when the radio MS radio interface is active and the whole simulation time. The radio interface is considered to be active if either not all the power saving classes of the MS are active, or each active class is in a listening window.

A. Interactive traffic analysis

To evaluate the performance of interactive applications we select a network scenario with only one MS with instant messaging traffic. In fact, instant messaging traffic has very low bit-rate, hence the interference with other MSs with the same type of traffic is very limited, and we ignore it so as to focus specifically on the performance with respect to the power saving parameters. Instant messaging is inherently symmetric and bi-directional, thus two connections are established, one in downlink and another in uplink, with the same statistical characterization. We assume that the size of each message is equal to the size of a single, i.e. non-concatenated, Short Message System (SMS), which is 140 bytes [15]. The arrival process of the messages is modeled via the Telnet ns-2 application [11].

We begin our analysis by evaluating the impact on the duty cycle of the following power saving parameters: T^{idle}, W^i, and W^l. To this aim, we vary each of them in the range of interest reported in Table 2, while the others are kept constant. For

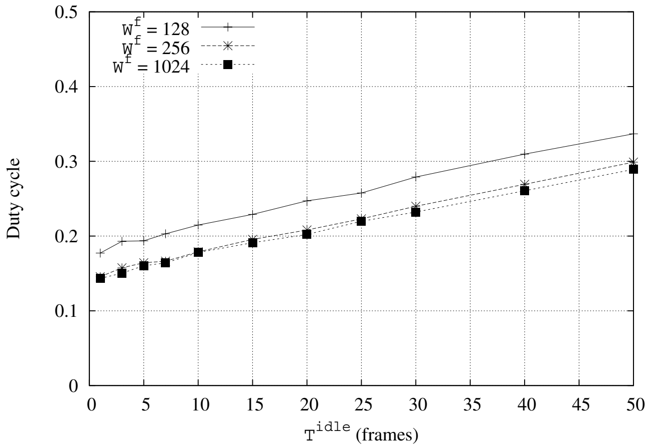

Fig. 2. Duty cycle vs. T^{idle}.

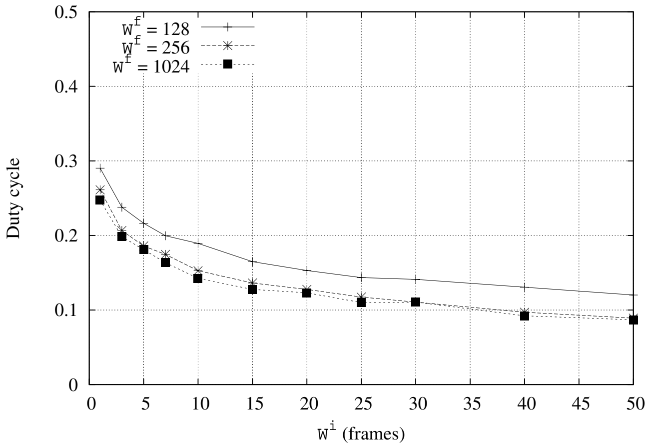

Fig. 3. Duty cycle vs. W^i.

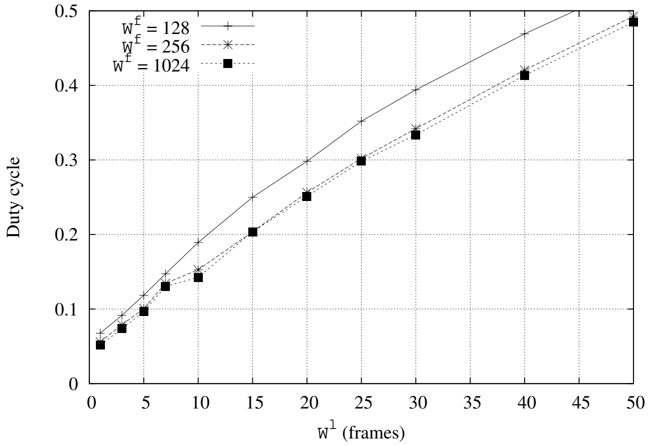

Fig. 4. Duty cycle vs. W^l.

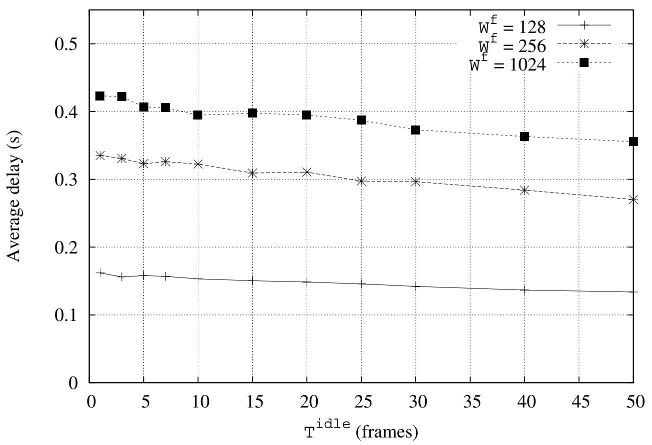

Fig. 5. Average delay vs. T^{idle}.

each case, we consider all the possible values of the W^f in Table 2, i.e. 128, 256, and 1024.

First, we show in Fig. 2 the duty cycle against the T^{idle} parameter, with $W^i = 10$ and $W^l = 2$. Since the responsiveness of activating a power saving class decreases when the idle time increases, all the curves in Fig. 2 increase. For all values of T^{idle} the curve with W^f equal to 128 lies significantly above the others, and the offset is almost constant. This is because W^f only affects the number of transitions from listening to sleep while a class is active, while T^{idle} the delay to activate a class.

Fig. 3 shows the duty cycle when W^i increases from 1 frame to 50 frames, with $W^l = 2$ and $T^{idle} = 2$. Unlike the case with increasing T^{idle}, the duty cycle here decreases when the initial sleep window W^i increases. This is because the larger the initial window is, the smaller the number of listening windows, and hence the energy consumption, becomes. In this case too, $W^f = 128$ entails worst performance, while the curves for W^f equal to 256 and 1024, respectively, are almost overlapping. Note that the energy consumption reduction is greater for small values of W^i, i.e. in the range between 1 frame and 20 frames.

Finally we report in Fig. 4 the duty cycle with W^l increasing from 1 frame to 50 frames, with $W^i = 10$ and $T^{idle} = 2$. As can be seen, the duty cycle is affected significantly by the listening window size. Specifically, the larger W^l is, the smaller the amount of energy saved becomes. In fact, the duty cycle difference between the minimum and maximum values in the range for T^{idle} and W^i is about 0.2, while that for W^l is almost 0.5. This can be explained as follows. The T^{idle} parameter only affects the performance once per class activation. Basically, W^i determines how many listening/sleep transitions are needed before the maximum sleep window size is reached. Instead, the W^l parameter is used every time the class enters a listening window, which results in a greater impact on the performance than the others.

We now analyze the performance in terms of the average delay, in the same set of scenarios above. We only report the curves for the downlink connection, since downlink is much more affected by varying the power saving parameters than uplink. This is because, in our implementation of power saving, a class becomes inactive as soon as a new MAC SDU is received at the MAC layer of an MS. In this case the MS immediately sends a message to the BS to notify that it just became awake and there are data waiting for transmission in the buffer of a given connection. On the other hand, when a MAC SDU is received by the BS, it is not possible for it to notify immediately the MS about this event, unless the active class is in a listening window. The BS will then have to wait until the first opportunity to send a traffic indicator message, which

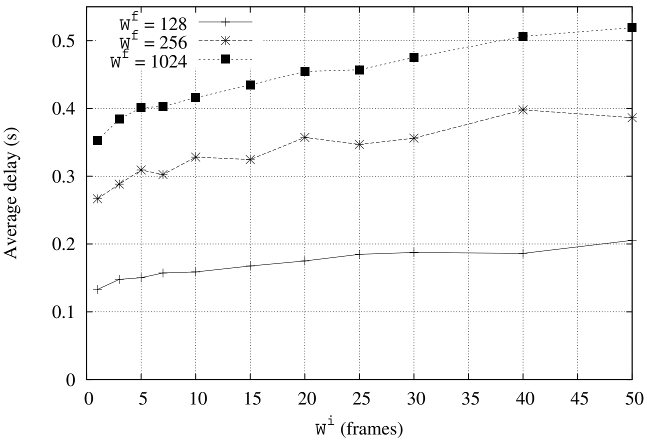

Figure 6. Average delay vs. W^i.

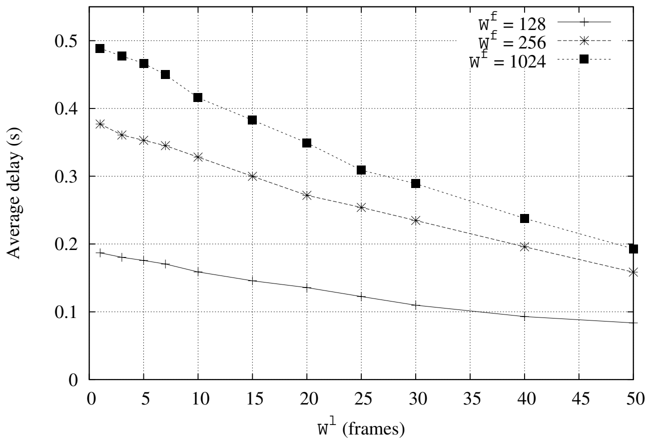

Figure 7. Average delay vs. W^l.

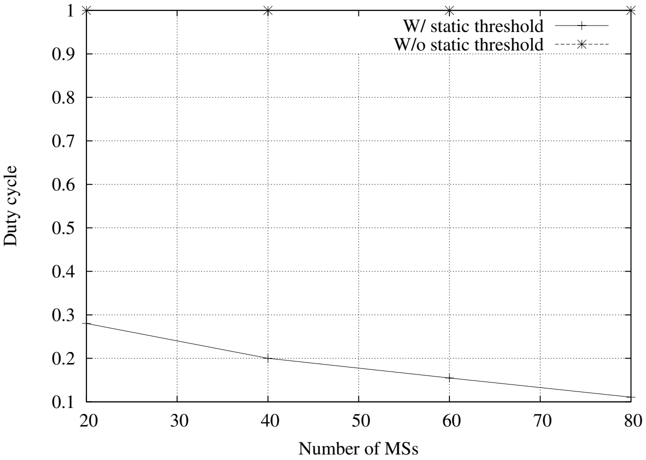

Figure 8. Duty cycle vs. number of MSs.

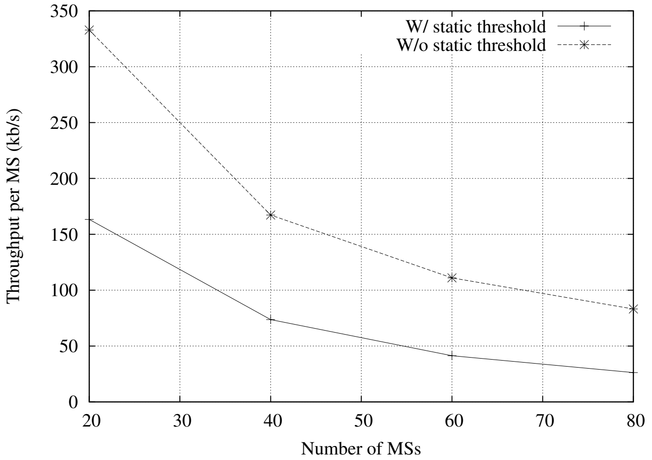

Figure 9. Throughput vs. number of MSs.

will entail the class de-activation.

As shown in Fig. 5, the average delay is only slightly affected by T^{idle}, which can thus be set so as to minimize the duty cycle (see Fig. 2), i.e. by setting T^{idle} as small as possible. Instead, the larger the final window size is, the greater the average delay becomes. Since values for W^f greater than 256 did not reduce significantly the amount of energy saved, we argue that such value yields a good trade-off between application performance and energy saving. The same behavior about the final window size is also verified in the rest of the scenarios discussed below.

We now show in Fig. 6 the average delay vs. the initial window size. As can be seen, there is about a two-fold increases of the average delay and of the duty cycle when W^i increases from the minimum, i.e. 1 frame, to the maximum, i.e. 50 frames. Since the duty cycle too is affected in a noticeable manner by setting this parameter, the optimal value depends on whether the system operator's policy is to value application performance over energy saving, or the other way around.

Roughly the same conclusion can be derived from Fig. 7, which shows the average delay when W^l increases from 1 frame to 50 frames. However, we note that in this case the duty cycle increases at a faster rate as W^l increases than the decreasing rate of the average delay. Therefore, setting W^l to a large value to improve the application performance should be considered carefully.

B. *Non-interactive traffic*

To evaluate non-interactive data traffic we employ a network scenario with an increasing number of MSs, from 20 to 80. Each MS has one TCP connection, whose 1024-byte segments are carried in the IEEE 802.16 network by means of a downlink connection (for data segments) and an uplink connection (for TCP acknowledgements). The *FTP* application of ns-2 has been used, and the TCP version was *NewReno*, with default parameters. The size of the per-connection buffer at the BS is 256 KB, which limits the number of bytes that the TCP sender can have in flight, i.e. with no acknowledgment from the TCP receiver.

Since non-interactive traffic, by definition, does not have specific time requirements, we implemented the following mechanism into the IEEE 802.16e BS MAC, called *static threshold*. When a MAC SDU is buffered, the connection is considered eligible for scheduling by the BS only if the total occupancy of the connection is greater than a threshold, in bytes. Otherwise, the connection is not granted any bandwidth. However, when data have been kept in the buffer for a duration greater than or equal to an expiration time then the connection is considered eligible for service immediately, regardless of the buffer occupancy. This is required to prevent service deadlock. Due to limited page budget, in the following we only report the results obtained with threshold equal to 256 KB and expiration time equal to 200 ms.

Figure 8 shows the duty cycle when the number of MSs increases from 20 to 80. As can be seen, without the static threshold mechanism, no energy is ever saved by the MSs. This is because the buffer at the BS always contains buffered data directed to all MSs. Thus, all requests to activate a power saving class from MSs are declined by the BS. Instead, by forcing the BS not to consider eligible those connections that have fresh new data buffered, i.e. by employing the static threshold mechanism, power consumption is greatly reduced. This result is consistent with the study in [10], where it is shown that with TCP traffic the RTT measured at the TCP sender significantly impacts on the performance, in terms of the amount of energy saved.

However, improving the performance in terms of energy has a side-effect: the throughput of the connection becomes lower. This is shown in Fig. 9, which reports the throughput against the number of MSs. In any case, the throughput decreases when the number of MSs increases, because the same amount of resources are shared by an increasing number of users. However, the curve obtained when static thresholds are enabled lies significantly below that without static threshold. Therefore, performance in terms of throughput has to be traded for energy saving.

IV. CONCLUSIONS

In this paper we studied the performance of power saving class type I of IEEE 802.16e via simulation. First, with interactive traffic, we varied the values of the following configuration parameters: listening window size, initial window size, final window size, idle time. The latter is not defined by the standard and has been introduced to determine when a connection is assumed to be idle after receiving a MAC SDU from upper layers. Results have shown that a final window size of 256 is a good trade-off between delay and energy saving. Also, the idle time should be set as low as possible since it affects only slightly the delay, but has a noticeable effect on energy saving. With regard to the listening window size and the initial window size, these parameters have a significant impact on both the delay and energy saving. Therefore, their tuning depends on how much delay the system operator is willing to introduce when at the same time improving the battery life of the mobile terminals.

Furthermore, we defined a mechanism to be implemented into the BS for elastic TCP traffic, called static threshold, which artificially increases the RTT. The above mechanism results in a significant gain in terms of energy saved, though lower throughput is achieved.

REFERENCES

[1] IEEE 802.16e, “IEEE draft standard for Local and Metropolitan Area Networks – Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems. Amendment for Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands”, Feb. 2007

[2] C. Cicconetti, C. Eklund, L. Lenzini, and E. Mingozzi, “Quality of Service support in IEEE 802.16 networks, IEEE Network, vol. 20, no. 2, pp. 50–55, 2006.

[3] G. Anastasi, M. Conti, E. Gregori, and A. Passarella, “A performance study of power saving policies for Wi-Fi hotspots”, Computer Networks (Elsevier), vol. 45, no. 3, pp. 295–318, 2004.

[4] G. Anastasi, M. Conti, E. Gregori, and A. Passarella, “Saving energy in Wi-Fi hotspots through 802.11 PSM: An analytical model”, proc. WiOpt 2004, Cambridge, UK, Mar. 24–26, 2004.

[5] R. Kravets and P. Krishnan, “Power management techniques for mobile communication”, proc. ACM/IEEE MOBICOM 1998, Dallas, TX, USA, Oct. 25–30, 1998.

[6] M.-G. Kim, M. Kang, J.-Y. Choi, “Performance evaluation of the sleep mode operation in the EEE 802.16e MAC,” proc. ICACT 2007, Phenix Park, Korea, Feb. 12, 2007.

[7] Y. Xiao, “Energy saving mechanism in the IEEE 802.16e wireless MAN,” IEEE Commun. Letters, vol. 9, no. 7, pp. 595–597, 2005.

[8] Y. Zhang and M. Fujise, “Energy management in the IEEE 802.16e MAC,” IEEE Commun. Letters, vol. 10, no. 4, pp. 311–313, 2006.

[9] M. E. Crovella and A. Bestavros, “Sel-similarity in world wide web raffic: evidence and possible causes,” IEEE/ACM Trans. on Networking, vol. 5, no. 6, pp. 835–846, 1997.

[10] J. Jang, K. Han, S. Choi, “Adaptive power saving strategies for IEEE 802.16e mobile Broadband Wireless Access,” proc. APCC 2006, Busan, Korea, Aug. 31–Sep. 2, 2006.

[11] Network simulator 2, http://nsnam.isi.edu/nsnam/

[12] A. M. Law and W. D. Kelton, “Simulation modeling and analysis,” Third edition, McGraw Hill, 2000.

[13] C. Cicconetti, E. Mingozzi, and G Stea, “An integrated framework for enabling effective data collection and statistical analysis with ns2,” Proc. WNS2 2006, Pisa, Italy, Oct. 10, 2006.

[14] WiMAX Forum, “Mobile system profile: Release 1.0 approved specifications,” May 2007.

[15] 3GPP TS03.40 Technical realization of the Short Message Service (SMS), v 7.5.0, 2002-02-11.