Computer Networks 57 (2013) 3276–3290

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

On the effectiveness of single and multiple base station sleep modes in cellular networks ☆

Marco Ajmone Marsan [a,b], Luca Chiaraviglio [c], Delia Ciullo [d,*], Michela Meo [a]

[a] Electronics and Telecommunications Department, Politecnico di Torino, Italy
[b] Institute IMDEA Networks, Madrid, Spain
[c] DIET Department, University of Rome La Sapienza, Italy
[d] Mobile Communications Department, EURECOM, Sophia-Antipolis, France

ARTICLE INFO

Article history:
Received 15 November 2012
Received in revised form 18 July 2013
Accepted 22 July 2013
Available online 30 July 2013

Keywords:
Energy-efficiency
Cellular networks
QoS
Sleep modes

ABSTRACT

In this paper we study base station sleep modes that, by reducing power consumption in periods of low traffic, improve the energy efficiency of cellular access networks. We assume that when some base stations enter sleep mode, radio coverage and service provisioning are provided by the base stations that remain active, so as to guarantee that service is available over the whole area at all times. This may be an optimistic assumption in the case of the sparse base station layouts typical of rural areas, but is, on the contrary, a realistic hypothesis for the dense layouts of urban areas, which consume most of the network energy.

We consider the possibility of either just one sleep mode scheme per day (bringing the network from a high-power, fully-operational configuration, to a low-power reduced configuration), or several sleep mode schemes per day, with progressively fewer active base stations. For both contexts, we develop a simple analytical framework to identify optimal base station sleep times as a function of the daily traffic pattern.

We start by considering homogeneous networks, in which all cells carry the same amount of traffic and cover areas of equal size. Considering both synthetic traffic patterns and real traffic traces, collected from cells of an operational network, we show that the energy saving achieved with base station sleep modes can be quite significant, the actual value strongly depending on the traffic pattern. Our results also show that most of the energy saving is already achieved with one sleep mode scheme per day. Some additional saving can be achieved with multiple sleep mode schemes, at the price of a significant increase in complexity.

We then consider heterogeneous networks in which cells with different coverage areas and different amounts of traffic coexist. In particular, we focus on the common case in which some micro-cells provide additional capacity in a region covered by an umbrella macro-cell, and we prove that the optimal scheduling of micro-cell sleep times is in increasing order of load, from the least loaded to the most loaded. This provides a valuable guideline for the scheduling of sleep modes (i.e., of low-power configurations) in complex heterogeneous networks.

1. Introduction

Eskimos are said to use many words for snow, because snow pervades their environment. On the contrary, in the early days of networking, the term *power* was used to

☆ Some of the results in this paper have been presented in [1,2].

* Corresponding author. Tel.: +33 0493008280.
E-mail addresses: ajmone@polito.it (M. Ajmone Marsan), luca.chiaraviglio@diet.uniroma1.it (L. Chiaraviglio), delia.ciullo@eurecom.fr (D. Ciullo), michela.meo@polito.it (M. Meo).

1389-1286/$ - see front matter
http://dx.doi.org/10.1016/j.comnet.2013.07.016

identify the ratio of throughput over delay [3,4], because energy issues did not belong in the networking landscape. Then, cellular networks and battery-operated terminals (most notably mobile phones) came along, and power control (now real power, measured in J/s) became an issue, in order to extend both the distance from the base station at which a terminal could be used, and the battery charge duration (in the early 1990s, heavy users carried a spare charged battery in their pocket, to avoid being cut off around noon). Next, sensor networks brought with them the question of power minimization to increase the network lifetime. Still, before the turn of the century, power consumption was not an element of the wired network design space. The first paper that addressed energy issues in fixed networks was [5], where Gupta and Singh investigated the energy consumption of Internet devices, and discussed the impact of sleep modes on network protocols. Since then, the interest in energy-efficient networking has been steadily rising, and the energy issue is now addressed in many conferences and research projects, among which we wish to mention TREND (Towards Real Energy-efficient Network Design) [6], the Network of Excellence funded by the European Commission within its 7th Framework Programme, which supported the work reported in this paper.

The directions that are presently pursued to achieve energy efficiency in networking can be grouped in two classes: (1) development of new technologies that reduce energy consumption and (2) identification of approaches that make the network energy consumption proportional to traffic. The rationale for the second direction derives from the observation that today network equipment exhibits power consumption which is practically independent of load. For example, a base station of a cellular network consumes at zero load about 60–80% of the energy consumption at full load [7].

Approaches that aim at improving the proportionality between the network energy consumption and the network load can be further divided in 2 sub-classes: (2a) development of equipment exhibiting better proportionality of energy consumption to load and (2b) identification of algorithms that allow the reduction of the functionality of network equipment in periods of low traffic, so as to decrease energy consumption in such periods. The algorithms that received most attention in class 2b are often called speed scaling, and sleep modes. By speed scaling we normally mean that the equipment can operate at different clock rates, with lower rates corresponding to lower power (and lower performance). By sleep modes we mean that in periods of low traffic the network operates with a subset of its equipment, the rest being switched off to save energy.

In the case of cellular networks, the critical equipment for power consumption is the base station (BS), whose typical consumption ranges between 0.5 kW and 2 kW [8,9], including power amplifiers, digital signal processors, feeders, and cooling system. Moreover, according to [10], all together, the BSs make up for about 80% of the total energy consumption of the cellular network.

In this paper we consider sleep modes for BSs in cellular networks, with reference to 3G technology, and we investigate the benefits that can be achieved by putting to sleep, i.e., bringing to a low-power-idle (LPI) state, a BS during periods of low traffic. This means that, in the future, the cellular access network planning should allow the selection of different operational layers corresponding to *network configurations* that specify the set of active BSs to serve different levels of traffic. These configurations can be activated according to predefined schedules, that are derived based on a combination of traffic forecasts and logs of traffic measurements. We compute the maximum amount of energy that can be saved with this approach, and we study the impact of the number of configurations, considering different types of network topologies with idealized cell structures. We then consider a real BS layout in a urban environment and a realistic coverage map, and show that significant savings can be achieved also in this scenario.

We assume that when some BSs are in sleep mode, radio coverage and service provisioning are taken care of by the base stations that remain active, so as to guarantee that service is available over the whole area at all times. This may be an optimistic assumption in the case of sparse base station layouts in rural areas (where network planning usually aims at coverage using large cells), but is on the contrary a realistic hypothesis for the dense layouts of urban areas (where network planning normally aims at capacity, with very redundant coverage, based on few large and many small cells), which consume most of the network energy. When some BSs enter the LPI state, the base stations that remain active may need to increase their transmission power, so as to cover also the area that was covered by the sleeping BSs. However, in our previous study [18] we showed that this increment in power consumption is usually negligible.

The main contributions of this paper are the following. We develop an analytical framework to identify the optimal scheduling of low-power network configurations (including how many BSs should be put into sleep mode and when) as a function of the daily traffic pattern, in the cases in which either just one low-power configuration per day is possible (bringing the network from a high-power, fully-operational configuration, to a low-power reduced-capacity configuration), or several low-power configurations per day are permitted (progressively reducing the number of active base stations, the network capacity, and the network power). We then compute the achievable energy savings in several cases: (i) assuming that any fraction of base stations can be put to sleep, (ii) accounting for the constraints resulting from typical regular base station layouts, and (iii) considering the case of a realistic network deployment in the city center of Munich. Moreover, we consider heterogeneous networks in which coverage is obtained by the superposition of macro-cells, that act as umbrella cells, and micro-cells, that provide additional capacity in specific areas. We prove that the optimal scheduling according to which micro-cells should be put to sleep is in order of increasing load. We show that, in a realistic heterogeneous network with real traffic profiles, large savings can be achieved by putting BSs to sleep, starting from the least loaded to the most loaded.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Optimal energy savings