Journal of Network and Computer Applications 116 (2018) 86–94

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

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

An optimal dynamic sleeping control policy for single base stations in green cellular networks

Sergio Herrería-Alonso*, Miguel Rodríguez-Pérez, Manuel Fernández-Veiga, Cándido López-García

Department of Telematics Engineering, University of Vigo, 36310 Vigo, Spain

ARTICLE INFO

Keywords:
Energy efficiency
Cellular networks
Base station sleep modes
Energy-delay trade-off

ABSTRACT

Energy efficiency of cellular networks can be greatly improve if base stations (BSs) can be put into a low power operation mode during low load periods. In this paper, we present a new dynamic scheme to manage the sleep mode of energy-aware BSs that can be autonomously governed. This kind of BSs, such as those operating in heterogeneous or hyper-cellular networks, can be put to sleep without causing undesired coverage holes, thus preserving the cellular service over the whole coverage area. With the proposed mechanism, energy-aware BSs enter the sleep mode as soon as they have no traffic to send and then apply a coalescing algorithm that dynamically adjusts the length of the sleeping periods to maximize the energy savings while maintaining the average service delay around a target value at the same time.

1. Introduction

Public awareness of climate change is driving current research efforts to augment energy efficiency in every technology field. In the telecommunications industry in particular, the ever increasing bandwidth demands and the expected increase in the number of connected devices make it dearly necessary to optimize the energy consumption of communication equipment, not only for the aforementioned environmental reasons, but also on plain economic grounds (Bolla et al., 2011; Fehske et al., 2011).

As all networks are designed with surplus capacity to accommodate traffic surges, and, as the actual demands change over time, a very successful technique to improve energy efficiency consists in powering down some parts of the network when they sit idle (Gupta and Singh, 2003). The end result is that energy profiles of communication equipment become more and more linear with traffic load as more opportunities to sleep are taken at the expense of some negative impact in the quality of service due to the diminished network capacity (Chiaraviglio et al., 2013). In mobile networks these techniques have already been successfully used at the device level, putting into a sleep mode unused resources in user equipment (UE) as, for instance, with the DRX protocol (Bontu and Illidge, 2009). However, about 60–80% of the energy required to operate a cellular network is consumed at the base stations (BSs) (Fehske et al., 2011), so improving their energy efficiency is critical to significantly reduce the overall power consumption of these communication systems.

The idea of putting redundant BSs into a sleep mode temporarily is not new (Marsan et al., 2013; Wu et al., 2015). In order to power down a BS, the traffic of its associated devices must be diverted to other nearby BSs in the same area with available capacity (Marsan et al., 2013; Niu et al., 2010; Conte et al., 2011; Oh et al., 2013). This usually requires that BSs cooperate with each other (or with a central controller) to preserve the cellular service over the whole coverage area. However, in some dense layouts, BSs can themselves decide whether to sleep without causing coverage problems. This is the case both in heterogeneous networks (HetNets) and in hyper-cellular networks (HCNs) (Wang and Zhang, 2014; Zhou et al., 2016). In HetNets, a macro cell is overlaid by a layer of small energy-efficient cells that can be powered down independently. In HCNs, control traffic is decoupled from data traffic so, when there is no data traffic, those BSs responsible of carrying traffic can be put to sleep while UEs remain associated to another BS that manages their control traffic. In this paper, we focus on these single cell scenarios in which a BS (a small cell in a HetNet or a data BS in a HCN) can be governed and put to sleep autonomously without rising an undesired coverage hole.

Any mechanism proposed to govern energy-aware BSs must decide when they should enter and exit the low power mode. Some sleep schemes wait until the BS has no more traffic to send and then, after a short waiting time (known as closed-down or hysteresis time), enter the sleep mode in the absence of any new task to serve (Niu et al., 2015;

* Corresponding author.
E-mail address: sha@det.uvigo.es (S. Herrería-Alonso).

https://doi.org/10.1016/j.jnca.2018.05.014
Received 13 September 2017; Received in revised form 16 March 2018; Accepted 9 May 2018
Available online 12 May 2018
1084-8045/

Guo et al., 2016). The hysteresis time tries to prevent frequent changes between the normal and the sleep modes, as transitions increase service delay and also consume energy. However, the energy consumed during the hysteresis time is merely wasted, so, we build here on the BS model developed in (Guo et al., 2016) to analyze the convenience of waiting before entering the sleep mode and propose a scheme in which the hysteresis time can be fully eliminated. With our proposal, BSs simply enter the sleeping state as soon as the last task in a busy period is served or remain active indefinitely if traffic load is excessively high.

There also exist different waking up strategies, namely waking up after a fixed predefined time, concatenating several sleeping periods until there is at least one task pending service, or just remaining asleep until the amount of tasks waiting for service reaches a certain threshold (Niu et al., 2015; Guo et al., 2016; Wu et al., 2013). In this paper, we propose a new and practical coalescing scheme based on the latter strategy that dynamically modulates the length of the sleeping periods according to the desired average delay and actual traffic conditions. We show that, for all but the most stringent target delays and highest traffic loads, BSs using our proposal always enter the sleeping state as soon as they have no traffic to send, thus maximizing energy savings.

The rest of this paper is organized as follows. Section 2 presents the related work. In Sect. 3 we provide a brief description of our proposal to then derive some of its main properties analytically. Section 4 presents our adaptive coalescing improvement. The proposal is evaluated through simulation in Sect. 5. Finally, we present our conclusions and future work in Sect. 6.

2. Background

Energy-aware BSs may support one or several low power operation modes in which some of their components can be switched off to save energy. In this paper, we concentrate on those BSs that can only make use of two operation modes: the normal active mode and a sleep mode in which they just consume a small fraction of the power needed for normal operation but they cannot serve any task at all. Any mechanism proposed to govern these energy-aware BSs must answer the two following questions:

1. When to enter the sleep mode (the sleep scheme).
2. When to exit the sleep mode (the wake-up scheme).

2.1. The sleep scheme

Obviously, energy-aware BSs should only be put to sleep when they are unused or under-utilized. If the sleep mode is triggered while some tasks are being served, nearby BSs must absorb the load of the BS being switched off. In this case, the transition from active to sleep must be designed in such a way that any interruption of ongoing tasks is prevented (Conte et al., 2011). Clearly, this requires some kind of cooperation among the BSs and may degrade the service. Therefore, we assume that autonomously operated BSs can only sleep when no task is being served.

To reduce the number of mode transitions, several papers propose to delay the transition to the sleep mode until the cell has been empty for a certain period (Niu et al., 2015; Guo et al., 2016). Thus, when the BS becomes empty, it waits for a while (the closed-down or the hysteresis time) before entering the sleep mode. If a new task arrives at the BS during this hysteresis time, the BS begins to serve it at once. The reasoning for remaining active is that, at the higher traffic loads, tasks inter-arrival times are so short that it is not worth to enter the sleeping state because it will not save enough energy to compensate for the energy consumed in the transitions. In return, if the BS eventually enters the sleep mode, the energy consumed while waiting is simply squandered, so it is important to properly tune the hysteresis time to maximize energy savings. In this paper, we prove that the hysteresis time can be fully eliminated under certain conditions and that, therefore, BSs should simply enter the sleeping state as soon as the last task in a busy period is served or remain active if the load is excessively high.

2.2. The wake-up scheme

A sleeping BS should wake up when existent traffic conditions imply a high risk of overload in nearby active BSs or an unacceptable quality of service for end users. The network status can be monitored by the sleeping BS itself, by neighboring active BSs or by the core network (Ashraf et al., 2011; Zhang et al., 2015). In the former case, some radio frequency components of the BS are required to work even in the sleep mode, which inevitably consume a certain amount of power. In the latter cases, since the sleeping BSs do not require to be aware of the traffic conditions, their sleep mode can be more energy efficient but their backhaul interface must be kept active to receive the wake-up request from the neighboring cells or the management system.

The main challenge of the wake-up scheme resides in choosing a good threshold to wake up the BS, neither too early, to avoid wasting energy unnecessarily, nor too late, to avoid poor performance (Niu et al., 2015; Guo et al., 2016; Wu et al., 2013). The most straightforward scheme just wakes up the sleeping BS after a fixed predefined time (single sleep scheme). This scheme can be easily improved if, at the end of each sleeping period, the BS checks the network status so that several sleeping periods can be concatenated while there is no task waiting for service (multiple sleep scheme). Alternately, some promising schemes maintain the BS into the sleep mode until the number of pending tasks reaches a certain threshold. In this paper, we focus on these latter coalescing schemes since, as we will show later, they may potentially maximize energy savings in a simpler way.

3. About the sleeping control policy

We start considering an energy-aware BS that employs a hysteresis sleep scheme to reduce the number of mode transitions. Thus, when the cell becomes empty, the BS will wait for D seconds before entering the sleep mode and provided that no task arrives during the hysteresis time, the BS will eventually be put to sleep. We assume that the transition time required to enter the sleep mode is short enough to be ignored in the analysis.

Regarding the wake-up scheme, we choose a coalescing technique that makes the BS wait to accumulate some tasks in the system before exiting the sleep mode. Specifically, a sleeping BS will just resume its normal operation when there are N tasks waiting for service. This coalescing technique can be easily applied to HCNs since in this kind of cellular networks control BSs can force users to employ a particular data BS for their traffic transmissions. However, its deployment in conventional cellular networks for uplink traffic is more involved, as current UEs facing a sleeping BS are typically served by the macro-BS covering the area, or by a neighboring active BS. To deploy this scheme in this scenario we could make use of the eICIC (enhanced InterCell Interference Coordination) feature standardized in the LTE-A protocol. With eICIC (also known as blanking), a BS mutes its data transmissions in (part of) a frame in order to not increase interference on nearby UEs associated to other BSs, but it still keeps transmitting its control information at a very low power, thus preventing that current UEs associated to the blanking BS switch to a different BS.

Finally, recall that a transition period (the setup time) is required to awake the BS and return to the active state. We assume that the setup time is in the order of a few seconds. Although this condition limits the number of the BS components that can be turned off and, therefore, the energy efficiency of the sleep mode, it allows to design a control scheme that works on short timescales.