Albrecht J. Fehske

Towards Reducing Energy Consumption in Mobile Access Networks

Beiträge aus der Informationstechnik

Mobile Nachrichtenübertragung
Nr. 64

Albrecht J. Fehske

Towards Reducing Energy Consumption in Mobile Access Networks

Flow Level Modeling & Deployment Strategies

VOGT

Dresden 2014

Bibliografische Information der Deutschen Bibliothek
Die Deutsche Bibliothek verzeichnet diese Publikation in der Deutschen Nationalbibliografie; detaillierte bibliografische Daten sind im Internet über http://dnb.ddb.de abrufbar.

Bibliographic Information published by the Deutsche Bibliothek
The Deutsche Bibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibiograpic data is available in the internet at http://dnb.ddb.de.

Die vorliegende Arbeit wurde unter dem Titel „Towards Reducing Energy Consumption in Mobile Access Networks“ als Dissertation an der Fakultät Elektrotechnik und Informationstechnik der Technischen Universität Dresden genehmigt.

Promotionskommission:

Vorsitzender: Herr Prof. Frank Ellinger (TU Dresden)
1. Gutachter: Herr Prof. Gerhard Fettweis (TU Dresden)
2. Gutachter: Herr Prof. Hamid Aghvami (King’s College London)
3. Gutachter: Herr Prof. Eduard Jorswieck (TU Dresden)

Tag der Einreichung: 14. Oktober 2013
Tag der Verteidigung: 10. Januar 2014

© Jörg Vogt Verlag 2014
Alle Rechte vorbehalten. All rights reserved.

Gesetzt vom Autor

ISBN 978-3-938860-72-4

Jörg Vogt Verlag
Niederwaldstr. 36
01277 Dresden
Germany

Phone: +49-(0)351-31403921
Telefax: +49-(0)351-31403918
e-mail: info@vogtverlag.de
Internet : www.vogtverlag.de

Abstract

The increasing demand for wireless services and ubiquitous access comes to the mobile communications industry at the cost of a sizeable carbon footprint as well as a significant energy bill. A considerable amount of the total energy of network operators is used to run the radio access network, i. e., the base stations. This thesis develops theoretical foundations to understand, assess, and minimize energy consumption in mobile access networks.

We first introduce a queueing-theoretic framework to analyze the load-dependent dynamic part of the network energy. Among others, low-complexity approximation techniques for the so-called coupled processor queue (which is analytically intractable) are developed. Such queueing-theoretic models are fundamental for a numerically efficient system analysis in general, since they allow to compute a variety of performance indicators (e. g., power) for a variety of deployments and traffic conditions without resorting to elaborate dynamic system level simulations.

We further characterize the relationship between network energy consumption per unit area and site densities, level of sectorization, as well as carrier frequencies for given coverage targets. Furthermore, we compare the average power requirements for two relevant strategies to expand existing networks: On the one hand, the deployment of low power nodes, called small cells, alongside conventional macro base stations, and, on the other hand, the installation of additional sectors at existing macro sites. The former strategy is heralded by network vendors and operators as a means to greatly increase network capacity and, at the same time, to yield significant energy savings compared to the latter. Our numerical results indicate, that small cell deployments lead to appreciable energy savings only in case of very high overall traffic demand or in case of very large differences between average and hotspot traffic intensity. In many cases, conventional deployments yield comparable energy consumption.

In a last part, we extend existing approaches to optimize user association policies, i. e., the rules that identify locations in the network with individual cells. Such techniques are of practial importance to adapt cell areas (in particular small cells) to non-uniform traffic distributions, which improves the quality of service perceived by the users. Application of optimized user association policies is shown to have only small effects on the base station power consumption. We argue, however, that these techniques are key to sustain quality of service in scenarios where cells are deliberately switched off to save energy.

Contents

1. Introduction **1**

2. Dynamic System Modeling **7**
2.1. Base Station Energy Consumption 7
2.2. Dynamic Traffic Demand 13
2.3. Network Layout, Radio Link, and Ressource Sharing 15
2.4. Flow Level Dynamics of Interfering Cellular Base Stations 20
2.5. Numerical Studies 38
2.6. Chapter Summary 42

3. Energy Aspects of Deployment Strategies **47**
3.1. Area Coverage and Energy Consumption 47
3.2. Capacity, Quality of Service, and Energy Consumption 56
3.3. Chapter Summary 74

4. Load-Aware User Association Policies **77**
4.1. Theoretical Framework 78
4.2. Numerical Studies 82
4.3. Chapter Summary 87

5. Summary And Conclusions **91**

A. Proofs **95**
A.1. Proof of Proposition 2.1 95
A.2. Proof of Proposition 2.2 97
A.3. Proof of Proposition 4.2 99

B. Simulation Model Parameters **101**
B.1. Propagation Models And Link Budget 101
B.2. Link Model Parameters 104
B.3. Base Station Power Model Parameters And Values 104

Acronyms

BS base station

KPI key performance indicator

SotA state-of-the-art

SINR signal-to-interference-and-noise-ratio

MIMO multiple-input-multiple-output

ISD inter-site distance

TTI transmission time interval

DTX discontinuous transmission

LTE Long Term Evolution

QoS quality of service

LOS line-of-sight

NLOS non-line-of-sight

BSD base station density

SFBC space-frequency block coding

HPBW half power beamwidth

PA power amplifier

RRH remote radio head

ICT information and communication technology

RAN radio access network

TCP transmission control protocol

IP internet protocol

1. Introduction

Rarely have technical innovations changed everyday life as fast and profoundly as the pervasive use of personal mobile communications. Over the last two decades, mobile wireless services grew from being niche market applications to globally available components of daily life. Today the total number of mobile subscriptions in the world has surpassed 4 billion, more than half the population of the planet. By comparison, there are only about 1.3 billion fixed line subscribers worldwide. Till date, the number of people accessing the internet amounts to only 2.4 billion worldwide which is about 35 % of the global population [1].

The increasing demand for wireless services and ubiquitous access, comes to the mobile communications industry at the cost of a sizeable carbon footprint. The entire information and communication technology (ICT) sector is estimated to represent about 2 % of global CO_2 emissions, a fraction comparable to that of global aviation [2]. According to a recent study, mobile networks accounted for 0.2 % of global CO_2 equivalent emissions in 2007 and are predicted to reach 0.4 % by 2020. Moreover, the ICT footprint is likely to less than double between 2007 and 2020, whereas the carbon footprint of mobile communications might almost triple within the same period [3].

The economic consequences resulting from energy consumption of mobile networks appear to be even more pressing than the ecological ones. In fact, the energy cost of running a mobile network can be significant. As a good rule of thumb, in mature western markets, the radio access network (RAN), i.e., the base stations (BSs), consume about 70 to 80 % of the total operational energy used by network operators and are accountable for about 30 % of their entire footprint. The corresponding costs constitute more than 20 % of overall OPEX - often times more than the personnel cost needed to operate the network [3]. In light of flattening or even decreasing revenues per user in mature markets, cost reduction becomes more and more important for mobile operators. Here, understanding and minimizing RAN energy consumption without compromising the quality of service (QoS) experienced by the customers is of key importance.

In developing countries, the situation is even worse. Due to a widespread lack of an electrical grid, many sites are powered by diesel generators today. On average, the transport of diesel to remote sites alone doubles the cost per liter compared to the price at the pump. In addition, costly power backup solutions are required for operation. As a result, conventional business models often times render infeasible in such markets, which is part of the reason why, as of today, there are still about one billion people worldwide that do not have access to any telecommunication services [4]. Operating BS sites with alternative energy, e. g., using locally installed solar panels or wind generators is part of the solution. The feasibility of these altenative energy sources, however, strongly depends on the average energy consumption of the equipment: Taking, for instance, solar panels with a peak output of 200 W/m^2 and 5 h daily peak sunlight, then a sizeable area of more than 20 m^2 of solar panels is required to power a threefold sectorized site with 40 W output power per sector and an average utilization of 40 %[1]. Clearly, reducing the average energy needs via intelligent deployment and management techniques will greatly facilitate integration of alternative power sources into base stations and thus open up possibilities for provision of wireless broadband access globally.

This thesis develops theoretical foundations to understand, assess, and minimize energy consumption in cellular networks with a focus on energy aspects of network deployment strategies. The following two sections narrow down the topics addressed in more detail.

Time Scales of Network Operation and Focus of this Thesis

Understanding its overall energy consumption requires taking a global view of a cellular network, which entails considering the complete system rather than individual, isolated components. Since there is no model that can capture all effects single-handedly, it is intuitive and practical to look at networks on different time scales, corresponding to the time scales on which different components operate, or changes in configuration are applied. Here, we consider three time scales, corresponding to the orders of milliseconds to seconds, minutes to hours, and days to weeks. The techniques and algorithms that may be applied on the smallest scale of seconds and milliseconds target mainly radio resource allocation and adapt individual transmissions to the wireless channel as well as to instantaneous load conditions. On the medium scale of minutes to hours, network management techniques adapt the network to the daily variations of traffic demands, e. g., during peak and off-peak hours. Finally, on the order of days to weeks, deployment strategies are applied to roll-out or expand a network according to coverage and peak traffic demands. Clearly, strategies on all time scales are interdependent and determine the overall network power consumption, which is further illustrated in Fig. 1.1.

The work presented in Chapter 3 and Chapter 4 of this thesis explicitely considers the network on the two larger time scales. In order to perform these investigations, however, we first require a model that captures the dynamics of the small time scale in detail and thus allows us to accurately estimate the average behavior observed on the larger scales.

[1]The underlying model of a site's power consumption is presented and discussed in more detail in Section 2.1.

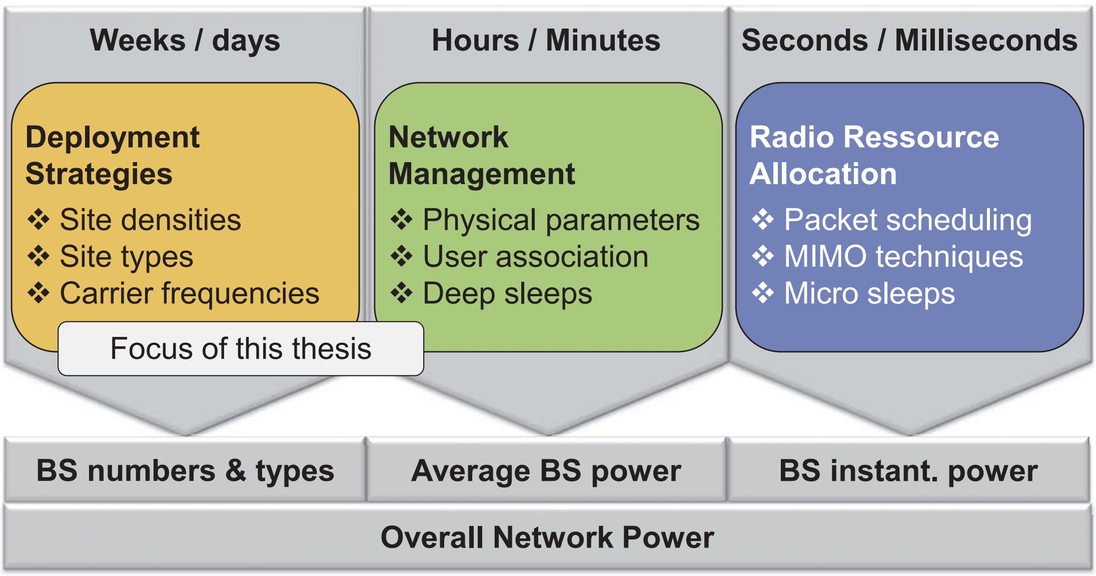

Fig. 1.1.: RAN optimzation and energy consumption on different time scales and focus of this thesis.

This model is developed in Chapter 2. The next paragraph outlines the contributions of this thesis in greater detail.

Outline and Contributions

This thesis aims at providing the theoretical basis for understanding, assessing, and minimizing the energy consumption in cellular radio access networks. The individual contributions are outlined as follows:

Network Modeling Framework

In Chapter 2, we provide a framework to analyze the dynamics of mobile communication networks. We are specifically concerned with queueing theoretic apporaches to assess the average BS resource utilization in the presence of inter-cell interference. The adequate queueing model, called the *coupled processor model*, renders analytically intractable and, till date, only very few analytical results exist to approximate the behavior of such systems. We extend current state-of-the-art (SotA) by proposing two techniques that closely approximate the dynamic system behavior. Such models are not only required to correctly estimate the BS powers, but also to assess other relevant key performance indicators (KPIs) such as throughputs and delays related to the QoS experienced by the users. Numerical studies undertaken in Chapter 2 suggest, that the techniques proposed provide closer estimates of the average utilization, the BS power consumption, and other KPIs than techniques known till date.

One of the approximation techniques proposed in this thesis is independently suggested in [5] within a different mathematical framework. The approach taken here provides two key advantages: Firstly, the approximation technique is seamlessly integrated into an overall queueing theoretic framework and compared to alternative techniques produced by this framework. Secondly, the method proposed here to numerically compute the BS loads, namely a fixed point algorithm, is conceptually and numerically significantly less involved than the method suggested in [5].

It is also noteworthy, that the applications of these approximate coupled processor models go beyond cellular networks. One of the techniques has already been succesfully applied to analytically estimate the performance of routers in many-core processor networks in [6]. In this application area, interfering BSs correspond to packet queues, which contend for access to ports at the routers.

Energy Consumption and Deployment Strategies

In Chapter 3, we utilize the framework proposed in Chapter 2 to characterize the relationship between deployment strategies and network energy consumption. Initially, we examine the relationship between a network's area coverage and its energy consumption. Here, besides site density and level of sectorization, we also study the impact of different carrier frequencies on the energy consumption for given coverage targets. Furthermore, we study the relationdhip between traffic demand and power consumption for selected deployments. We compare, specifically, the average power for two relevant strategies to expand existing networks: On the one hand, deployment of low power nodes, also called *small cells*, alongside conventional macro BSs, and, on the other hand, setup of additional sectors at existing sites. Increasing spatial frequency reuse through small cells is heralded by operators and equipment vendors [7, 8] as a means to greatly increase network capacity and cope with strongly inceasing traffic demands.

Varying traffic demands imply varying levels of BS utilization, QoS, as well as power consumption, which need to be evaluated. The framework derived in Chapter 2 is fundamental for these studies, since it allows us to compute power consumptions for a variety of deployments without the need to resort to dynamic system simulations.

Load-Aware User Association Policies

Last but not the least, Chapter 4 deals with optimization of user association policies, which are the rules that associate any location in the network with a serving BS. Flexible user association is of great importance especially in small cell scenarios: Due to the limited range of these low power nodes and practical restrictions on their placement, cell areas do not necessarily match traffic hotspots in practice. So called cell range expansion techniques are required to ensure, that sufficent demand is actually covered by small cells. The corresponding theoretical framework in Chapter 4 is based on results presented in [9], which we extend with the models developed in Chapter 2 to accomodate the effects of load coupling through interference.

The investigations in Chapter 3 assume perfect alignment of the positions of both traffic hotspots and small cells, which consequently leads to upper bounds on the values of a variety of KPIs, due to displacement in practice. Numerical results in Chapter 4 indicate, that these values are attainable, at least for moderate displacements, by applying optimized user association policies.