

A Distributed Power Saving Algorithm for Cellular Networks

Ingo Viering¹, Matti Peltomäki², Olav Tirkkonen^{3,4}, Mikko Alava²,
and Richard Waldhauser⁵

¹ Nomor Research GmbH, Munich, Germany

² Department of Applied Physics, Helsinki University of Technology (TKK), Finland

³ Department of Communications and Networking, TKK, Finland

⁴ Nokia Research Center, Helsinki, Finland

⁵ Nokia Siemens Networks, Munich, Germany

Abstract. We consider power saving in a cellular network. Subject to constraints generated by network planning and dynamic Radio Resource Management algorithms, there is room for reducing Base Station (BS) transmit powers. We suggest that power is reduced in a way that does not move cell boundaries significantly. For this, there is a power difference constraint d_{\max} that upper limits the *difference* in power reductions in neighboring BSs. The BSs exchange information about their current level of power. Each BS has a target maximum power reduction. We propose a distributed algorithm, where the BSs take turns to reduce their power, respecting their neighbors current power settings and the power difference constraint. We show that this algorithm converges to a global optimum.

1 Introduction

Self-organization is a wide ranging research trend in modern networking. In the scope of wireless networking, research on Self-Organized Networks (SON) ranges from general principles of cognitive and ad hoc networks to concrete problems in standardization and implementation of near future mobile networks [1,2]. Here, we concentrate on a specific SON problem of current interest for the standardization of the next release of the Evolved Universal Terrestrial Radio Access Network (E-UTRAN), a.k.a. Long Term Evolution (LTE). The SON-related requirements for LTE-Advanced [3] are:

- energy efficiency of NW and terminals,
- SON in heterogeneous and mass deployments,
- avoiding drive tests.

This paper discusses a distributed algorithm for realizing energy savings in cellular networks. As such, it deals directly with energy efficiency, and with heterogeneous deployments—the only truly scalable way of implementing networking algorithms for heterogeneous mass deployments is to distribute the algorithms to the deployed network elements themselves.

The distributed algorithm for power saving is designed to respect cell boundaries up to a degree. It can be proven to converge in linear time to a global optimum.

2 Motivation for Power Saving

Power efficiency has not been an issue in network planning, and modern cellular communication networks such as LTE have been designed to operate in an interference limited manner. As a consequence adapting the transmit power will be an important SON feature. The target is to minimize output power without endangering coverage.

As an example we take the evaluation scenario “Case 1” defined in 3GPP [4]. This is a heavily interference limited scenario with the Inter-Site Distance 500m, and a Tx power of 46dBm. There is a significant potential for power saving. Figure 1 shows the average sector throughput in this scenario with a full buffer traffic model, and different Tx powers at the BSs. A reduction of 12dB (to 34dBm) leads to a negligible loss of 1.9% in throughput—the network remains interference limited. Based on this we argue that especially in small cells, there may be significant room for optimizing the transmit power of base stations.




Fig. 1. Loss in throughput when reducing base station transmit power

3 Distributed Network Power Saving

We propose an algorithm for reducing the power in cellular networks, where the base stations decide the reductions in the cells, based on local information such as load, and reference signal receive power (RSRP) measurements received from User Equipments (UEs).

As discussed in [5], an important requirement of power adaptation is that the coverage areas of neighboring cells shall not be affected, at least not by an individual BS. In other words, power adaptation of an individual cell shall not change the cell boundaries. The cell boundaries are carefully desinged in the network planning phase, and there will be other SON algorithms which may

dynamically optimize the boundaries for traffic reasons etc. To avoid clashes, we allow an individual cell to reduce its power only if all neighbors reduce their power in a rather similar way. For this, we assume a maximum power reduction difference d_{\max} , which is set by a central entity. An initial power difference may exist for two neighboring cells, designed in network planning and by other SON algorithms. For energy saving purposes, a cell may save power in a way that increases the power difference to neighbors with at most d_{\max} . As an example, assume that there are two cells, and in normal operation (i.e. given by network planning) the transmit power in cell A would be 46dBm, and in cell B 43dBm. In a power saving phase, cell A would like to reduce its power by 12dB to 34dBm and cell B by 21dB to 22dBm. If $d_{\max} = 2\text{dB}$, cell B could not reduce its power by more than 12dB+2dB, i.e. not to less than 29dBm.

3.1 System Model

The power reduction problem is defined as follows. There is a graph $G = G(V, E)$ of interference couplings between base stations, based on e.g. Neighbor Cell Lists. There are N BSs, which are the vertices of the graph. The set of neighbors of BS n is denoted by \mathcal{N}_n . Each node n uses initially a power $p_{n,0}$, and makes periodic decisions related to a power reduction r_n so that the transmit power used after the decision is $p_n = p_{n,0} - r_n$. Each base station has a maximum power reduction m_n such that under all circumstances $r_n \leq m_n$. This power reduction is selected so that cell coverage is not jeopardized—the minimum power base station n could use is $p_{n,0} - m_n$, otherwise it would generate a coverage hole. The power reductions of two adjacent base stations may not be more than d_{\max} units apart, or $|r_n - r_m| \leq d_{\max}$ for all $(n, m) \in E$. The task is to minimize $\sum_v p_n$ subject to the constraints above, or equivalently, to maximize $\sum_v r_n$. For convenience, we measure all powers, power reductions and differences in dB-scale.

In the distributed power reduction problem, the minimization is to be done such that (a) nodes decide independently how much power to use, (b) power reduction difference constraint is respected in all intermediate stages of the process, and (c) the global minimum is achieved.

3.2 Suggested Algorithm

The Distributed Power Saving (DiPoSa) algorithm works as follows:

1. There is a predetermined periodicity determining the moments when each BS makes a decision related to its power usage.
2. For simplicity, the algorithm is assumed to be iterative and asynchronous—the decision period is assumed to be the same for all BSs, and it is assumed that the BSs make decisions in an unsynchronized manner. This translates to the situation that the distributed algorithm sweeps through the BSs in a random order, and then sweeps through again in the same order. One sweep is called an iteration.¹

¹ Note that this is readily generalized to different periods for different BSs, or even periods changing from decision to decision.

3. At the moment BS n makes a decision related the power to use, it knows the power reductions $\mathcal{R}_n = \{r_m\}_{m \in \mathcal{N}_n}$ used by all its neighbors m at that moment.
4. The BS selects the power reduction $p_n = \min(\{m_n\} \cup \{r_m + d_{\max}\}_{m \in \mathcal{N}_n})$

In the last step the BS computes the maximum power reduction allowed by the powers used by each neighbour, by adding the power reduction of the neighbour and the maximum power reduction difference d_{\max} . If the smallest power reduction allowed by the neighbours is larger than the target maximum reduction m_n , the BS uses the smallest power reduction allowed by the neighbours. Otherwise, it uses the target maximum reduction.

3.3 System Requirements

DiPoSa is distributed so that the decision of the power used is made at the BS. In order to make the decisions, the BS needs to know the maximum power reduction d_{\max} , and the power reductions at use by its neighbors \mathcal{R}_n . The former is assumed to be signaled by Operations and Maintenance, and it is time-invariant. Note that it could even be different for different cells, but for the sake of illustration we have assumed that it is the same in the network. The latter are regularly exchanged directly between the BSs. In LTE, the network of BSs is meshed by direct BS-to-BS X2-interfaces, which is the natural candidate for exchanging \mathcal{R}_n information.

3.4 Convergence to Global Optimum

The objective function that each node minimizes is completely local, but it is connected by the power reduction constraint to the neighbors. In principle, connections may link up to generate long-distance connections between nodes. Nevertheless, the distributed algorithm converges to a global optimum, and in a time that is linearly proportional to $\max_n m_n/d_{\min}$.

Proposition 1. In a finite network, the DiPoSa algorithm described in Section 3.2 converges to a global optimum.

Proof: First we have to prove that DiPoSa converges to some state. This is easiest to do by observing that the BSs participate in an exact potential game [6]. There is a simple global potential function which is a sum of the power reductions achieved in the network. The utility function that each BS tries to maximize at each update is the individual power reduction. At each update, the increase in the updating nodes utility corresponds exactly to the increase in the global potential. Thus the settings for a exact potential game are fulfilled and DiPoSa converges. Next, we prove by contradiction that the converged state of DiPoSa is the global optimum. Assume that the algorithm has converged to a state $\mathcal{S} = \{r_n\}$, and that the global optimum corresponds to $\mathcal{O} = \{o_n\}$. If $\mathcal{S} \neq \mathcal{O}$, $o_n > r_n$ for at least a $n = n_1$. Due to Step 4 in DiPoSa, the reason for this must be that for at least for one neighbor of n_1 , say n_2 , we have $o_{n_2} > r_{n_2}$, and that

$$r_{n_1} = r_{n_2} + d_{\max} < o_{n_1} \quad (1)$$

Due to Step 4 in DiPoSa, a neighbor of n_2 has to be away from optimum. This neighbor of n_2 cannot be n_1 , as from (1) we have $r_{n_1} > r_{n_2}$. Thus there exists at least another neighbor of n_2 , say n_3 , for which $o_{n_3} > r_{n_3}$. Repeating this induction step N times we observe that the last BS is also away from optimum, $o_{n_N} > r_{n_N}$, and from equations induced from (1) we know that $r_{n_N} < r_n + d_{\max}$ for all $n \neq n_N$. Thus in Step 4 for n_N the reduction r_{n_N} can be increased. This contradicts the assumption that \mathcal{S} is a converged state. ■

A slightly more involved argument would show that DiPoSa converges in linear time, which is upper bounded by a linear function $t_{\text{conv}} \propto \frac{\max_n m_n}{d_{\max}}$.

4 An Example

Figure 2 gives an example for a heterogeneous cell layout with different cell sizes, modeling the cells covering the center and outskirts of a medium-size city. The network setting is generated as a Voronoi diagram, where the x- and y-coordinates of the underlying point set is Laplacian distributed. Every cell bears two numbers. The first contains the individual (uncoordinated) maximum power reduction m_n of every cell. These are determined in an ad hoc manner, the larger the cell the smaller m_n . The second number is the r_n in the converged state of



Fig. 2. Example network with maximum power reductions and reached optimum power reductions

DiPoSa, how much every cell is allowed to actually reduce the power. Here, we have used $d_{\max} = 2\text{dB}$. If the original transmit power is the same for all cells, i.e. $p_{n,0} = p_0$, the converged solution saves 67% of the overall output power in the example network. Without the d_{\max} constraint, i.e. if every BS n reduced by m_n , the savings would be 70%.

5 Conclusion

In this paper, we proposed a distributed power reduction algorithm, where each BS has a local target power reduction, and the reductions at use at the BSs should not differ between neighbors by more than a maximum d_{\max} . The objective functions for each BS are decoupled, but the constraint d_{\max} reduces the action space of the BSs, and induces long-range correlations. We observed that these long-range correlations do not prevent a distributed solution. The proposed algorithm is a straight forward implementation of the limitations of the local action spaces by which these long range correlations become realized in the converged solution. Due to the locality of the objective functions, the algorithm was proven to converge to the global optimum.

Acknowledgments. This research is partially funded by Nokia.

References

1. 3GPP: Self-configuring and self-optimizing network use cases and solutions. Technical Report TR 36.902 (2008), <http://www.3gpp.org>
2. Döttling, M., Viering, I.: Challenges in mobile network operation: towards self-optimizing networks. In: Proc. IEEE ICASSP (April 2009)
3. 3GPP: Requirements for further advancements for E-UTRA (LTE-Advanced). Technical Report TR 36.913 (2008), <http://www.3gpp.org>
4. 3GPP: Physical layer aspects for evolved universal terrestrial radio access (UTRA). Technical Report TR 25.814 (2009), <http://www.3gpp.org>
5. 3GPP: De-centralized optimization of downlink transmit power. Technical Report R3-091929 (2009), <http://www.3gpp.org>
6. Monderer, D., Shapley, L.: Potential games. *Games Econom. Behav.* 14, 124–143 (1996)