Universität Stuttgart

INSTITUT FÜR KOMMUNIKATIONSNETZE UND RECHNERSYSTEME
Prof. Dr.-Ing. Andreas Kirstädter

Institute of Communication Networks and Computer Engineering
University of Stuttgart
Pfaffenwaldring 47, D-70569 Stuttgart, Germany
Phone: ++49-711-685-68026, Fax: ++49-711-685-67983
Email: mail@ikr.uni-stuttgart.de, http://www.ikr.uni-stuttgart.de

Interference Coordination in Cellular OFDMA Networks

Marc C. Necker, University of Stuttgart

Abstract

Orthogonal frequency division multiple access is the basis for several emerging mobile communication systems. Prominent examples are the 3GPP Long Term Evolution as the successor of UMTS high-speed packet access and the IEEE 802.16 system, advanced by the WiMAX forum. On a system level, OFDMA is basically a combination of time and frequency division multiple access. In cellular TDM/FDM systems, inter-cell interference is a major issue that traditionally has been solved by avoiding the use of the same frequency bands in adjacent cells. However, this solution incurs a waste of precious frequency resources. An attractive alternative is the use of beamforming antennas in combination with interference coordination mechanisms, where the transmission of adjacent base stations is coordinated to minimize inter-cell interference. Interference coordination is an important aspect of the system level, which influences many other issues, such as network planning or scheduling mechanisms. In this article, we give an overview of interference coordination as it would apply, for example, to IEEE 802.16e and review the relevant literature. We also discuss and compare interference coordination algorithms, which can be based either on global system knowledge or purely on local system knowledge.

Orthogonal frequency division multiplexing (OFDM) is a widely deployed technology. It offers many advantages in multipath fading environments and has been used successfully for digital video broadcasting (DVB-T) and wireless local area networks (e.g., 802.11a). Several emerging standards for cellular broadband networks, such as 802.16e (WiMAX) or the future 3GPP Long Term Evolution (LTE) are based on orthogonal frequency division multiple access (OFDMA). In OFDMA, several users are multiplexed in time and frequency, based on an underlying OFDM system. A major problem in these systems is the inter-cell interference that neighboring cells cause when using the same frequency band, eventually leading to severe performance degradation or loss of connection.

Inter-cell interference can be mitigated in several ways. Code division multiple access (CDMA)-based systems, such as the universal mobile telecommunications system (UMTS) or cdma2000 use different scrambling codes in different cells, thus reducing inter-cell interference. On the other hand, classical frequency and time division multiplexing (FDM/TDM) systems, like global systems for mobile communications (GSM), avoid reusing the same set of frequencies in neighboring cells by employing a frequency reuse pattern. A third possibility is to apply space-division multiplexing (SDM) with beamforming antennas. By focusing their transmission or reception in the direction of a particular terminal, the interference to terminals in other directions is minimized. Finally, the transmissions of neighboring base stations can be coordinated further, thus almost completely eliminating inter-cell interference [1]. This is referred to as interference coordination (IFCO).

IFCO in combination with beamforming antennas is of interest to 3GPP LTE and 802.16e because it seems the most promising approach to solve the problem of inter-cell interference in OFDMA systems while at the same time achieving a high spectral efficiency. This article gives an overview of interference coordination in OFDMA systems in combination with beamforming antennas. In the following, we introduce elements of the 802.16e standard to illustrate the use of OFDMA in current standards. We then review the current literature on interference coordination and describe the most important relevant performance metrics for the study of interference-control algorithms in cellular networks. Finally, we discuss algorithms based on global system-state information and algorithms based purely on local-system state information. We present representative performance results (assuming otherwise ideal conditions) and comment on the capabilities and shortcomings of interference coordination based on local system information.

An Example OFDMA System — IEEE 802.16e

The IEEE 802.16e standard has a large number of configuration options. In addition to standard parameters like timer values, these comprise fundamental system properties like modulation and coding schemes. Therefore, the WiMAX forum has defined a system profile for operation and deployment of mobile IEEE 802.16e networks [2]. This profile fixes many of the various choices available in IEEE 802.16e to simplify product development and ensure product compatibility.

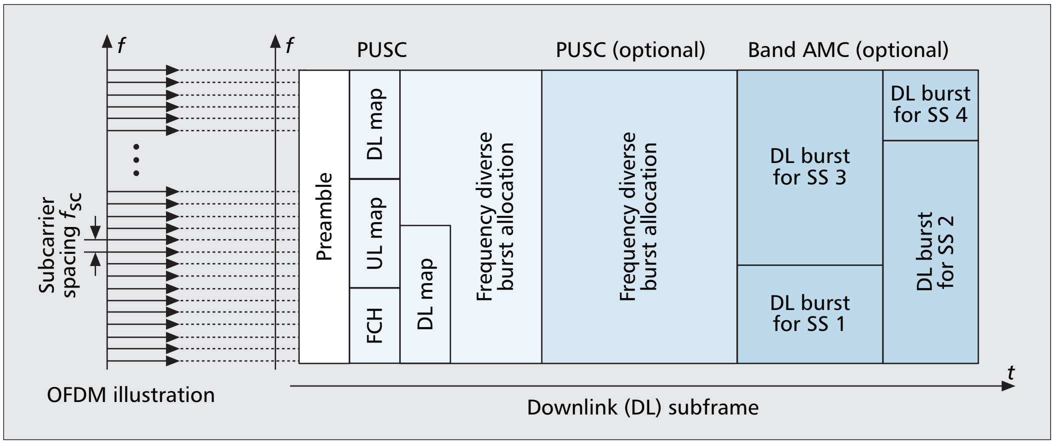

Figure 1. *Schematic view of an OFDM transmission system (left side) and the downlink portion of a TDD MAC frame (right side).*

Although IEEE 802.16e networks often are referred to as WiMAX networks, actually this name only should be used for networks conforming with the WiMAX Forum mobile system profile.

OFDM is a spread-spectrum transmission technique, which subdivides the available spectrum into a large number of orthogonal *subcarriers* with a carrier-spacing of f_{SC}. This is illustrated in Fig. 1. OFDM has many advantages, such as robustness in multipath fading environments, high spectral efficiency, and simple implementations by means of fast fourier transform (FFT). The overall system bandwidth can be scaled from 1.25 MHz up to 20 MHz in IEEE 802.16e, where the WiMAX mobile system profile specifies systems only up to a bandwidth of 10 MHz. The subcarriers of the underlying OFDM transmission scheme are grouped into *subchannels*. This can be done such that the subcarriers of a subchannel are spread across the frequency spectrum in a pseudo-random fashion, or it can be done by allocating contiguous subcarriers. These strategies aim at an exploitation of frequency diversity and frequency selectivity, respectively. The basic resource assignment granularity for data transmissions is a *slot*, which is defined by a subchannel index over a certain number of OFDM symbols. A slot always contains 48 data subcarriers plus additional pilot symbols. Several slots are grouped to form a *data burst*, within which one or more connections of a particular mobile terminal (also referred to as a subscriber station [SS]) is served.

In IEEE 802.16e time division duplex (TDD) systems, a MAC frame is divided into an uplink and a downlink subframe. Both subframes are further divided into permutation zones allowing for different operational modes. A sample downlink subframe is shown in Fig. 1. The first OFDM symbol of a frame forms the preamble that is used for synchronization and equalization. Subsequently, every frame begins with a mandatory partial usage of subcarriers (PUSC) zone containing control information and data bursts, followed by further optional PUSC zones. PUSC is one of the basic zone types operating in a frequency-diverse mode, where the subcarriers of a subchannel are spread across the frequency spectrum according to a predefined permutation scheme. In addition to PUSC, the full usage of subcarriers (FUSC) zone provides frequency diversity based on a different permutation scheme.

A third important zone type in addition to PUSC and FUSC is the band adaptive modulation and coding (band AMC) zone. In this zone, a set of contiguous subcarriers in frequency and OFDM symbols in time direction are allocated to one mobile terminal, that is, a frequency-selective resource allocation is performed. In the downlink, the standard requires the allocation of rectangular areas, as indicated in Fig. 1. In the uplink, a linear allocation of slots is foreseen. Similar allocation restrictions exist for the PUSC and FUSC zones, albeit with a less obvious geometric analogy.

IEEE 802.16e provides support for advanced antenna systems (AASs), both for beamforming and multiple-input multiple-output (MIMO). For beamforming systems, it is important that the receiver can still estimate the channel transfer function. Commonly, broadcast pilot symbols are used, and the mobile terminal uses all pilot symbols for channel estimation. Beamforming requires dedicated pilots, where every mobile terminal receives pilots specifically beamformed for itself. In the downlink, PUSC and band AMC both support dedicated pilots. In addition, fast channel quality feedback is supported to control the quality of downlink beamforming.

Interference Coordination in Ad Hoc and Cellular Networks

Classification

IFCO schemes are classified with respect to their time-scale of operation into *static schemes*, *semi-static schemes*, and *dynamic schemes*.

- *Static IFCO* is performed during the network planning process, for example, by an optimized frequency planning in the network. Such schemes operate in the time scale of days or longer.
- *Semi-static schemes* can be subdivided into three subclasses [3]. First, self-configured coordination schemes are basically a self-optimizing version of static schemes. Second, cell load adaptive coordination requires load estimates of the cells and operates on the level of minutes. Third, user-load adaptive coordination takes into account the instantaneous traffic demand, which varies faster than the cell load. Therefore, they operate on a time scale in the order of seconds or several hundred milliseconds. Metrics for the traffic demand, for example, can be the buffer occupancy or the quality of service (QoS) profile of a particular traffic flow.
- *Fully dynamic IFCO schemes* can adapt instantly to changing network conditions, such as changing traffic or load distributions. Their time scale of operation is in the order of one or only a few MAC frames.

With respect to the degree of distribution, four classes can be distinguished.

- *Global schemes* require an omniscient central entity that can capture the system state instantly and distribute scheduling decisions to all network nodes on a frame level. They

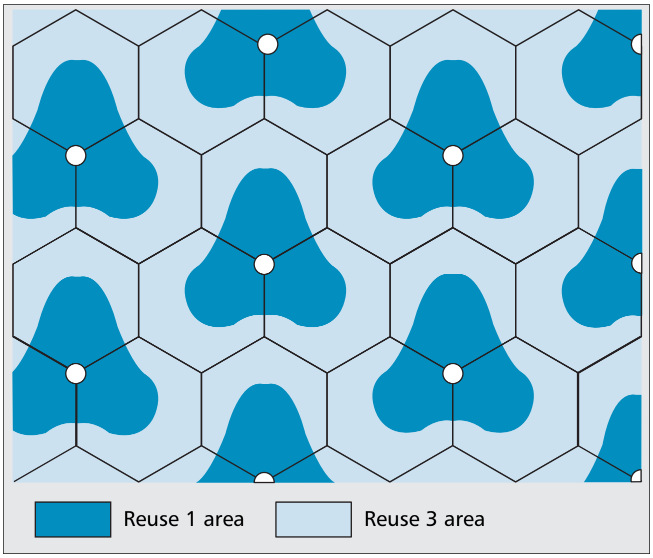

■ Figure 2. *Fractional frequency reuse.*

enable the study of basic system properties and trade-offs and may serve as a basis for the development of more practical schemes.

- *Distributed schemes* require a central entity but are designed to cope with real-world signaling delays and loads.
- *Decentralized schemes* operate without a central entity but allow base stations to exchange information.
- *Local schemes* do not require direct communication among base stations and operate only on local state information.

In the following, we first review certain important IFCO algorithms proposed in the literature. Subsequently, we illustrate the performance of a global and a local IFCO approach, respectively.

Selected Approaches from the Literature

In circuit-switched networks, techniques for dynamic channel assignment (DCA) were investigated as early as 1971 by Cox and Reudink [4]. Compared to classical frequency reuse with a static assignment of channels to cell sectors, DCA enables the dynamic reassignment of channels from one cell to another. A second basic scheme besides DCA is the borrowing channel assignment (BCA), which was proposed by Engel in [5]. In contrast to DCA where basically all cell sectors can use all channels, BCA begins with a fixed channel assignment and then allows a fully loaded cell sector to borrow unused channels from one of its neighbor sectors.

Another popular technique is reuse partitioning, which was introduced for circuit-switched networks by Halpern in 1983 [6]. It is a form of static planning and aims at increasing the capacity by using different frequency reuse factors for certain mobile terminals. Within the activities of the WiMAX forum and 3GPP LTE standardization, it is known as *fractional frequency reuse (FFR)*. Figure 2 illustrates FFR using the example of two different reuse factors for terminals close to the base station and terminals close to the cell edge. A generalization of FFR was studied by Bonald et al. in [7]. The authors define arbitrary geographic regions within a cell area, where in each region a terminal is served with a certain transmission profile. A transmission profile corresponds to a particular combination of active transceivers. Optimal region boundaries are then derived. The authors present numerical examples for a two-cell and a three-cell network and for infinitely large linear and hexagonal networks.

A dynamic local coordination scheme was proposed by Sternad *et al*. in [8]. They coordinate the scheduling among cell sectors of the same base station and mitigate the interference from neighboring base stations by additionally applying reuse partitioning. Altogether, the authors achieve an effective frequency reuse of 1.5–2 with an average signal to interference ratio (SIR) of 16dB.

IFCO also has been an active research area in multihop and mobile ad hoc networking for a long time. In [9], the authors consider the possibility of beamforming in a multihop wireless network and study a MAC protocol that is capable of blocking the transmissions of the strongest interferers. In [10], the authors coordinate broadcasts in a multihop wireless network by means of a sequential graph coloring heuristic. In [11], the coordination of transmissions in a wireless ad hoc network is considered. The interference conditions are evaluated by an omnipotent central entity with full system-state information, which is able to schedule the data transmissions of the individual nodes on the MAC-frame level. This is done based on a conflict graph, which represents critical interference relations in between the network nodes. The problem was traced back to the graph coloring problem, for example, in [12]. In [13], the throughput capacity of a wireless multihop network was calculated with the help of a very similar schedule graph, which is derived from the physical layer properties of the network.

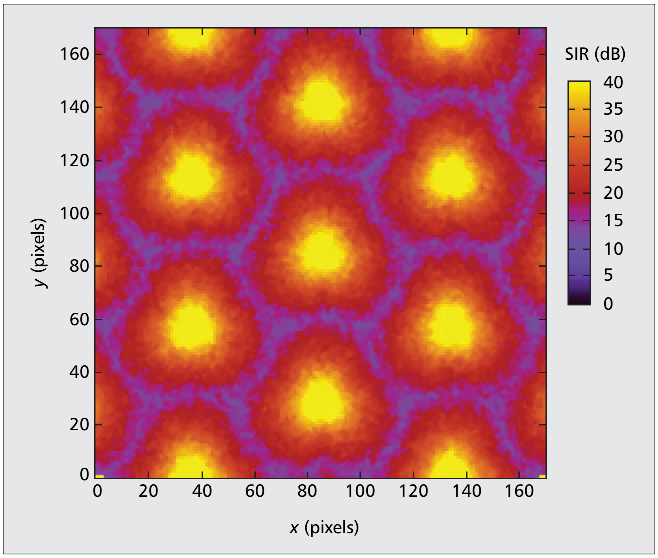

■ Figure 3. *Mean SIR, linear array antennas, frequency reuse 3.*

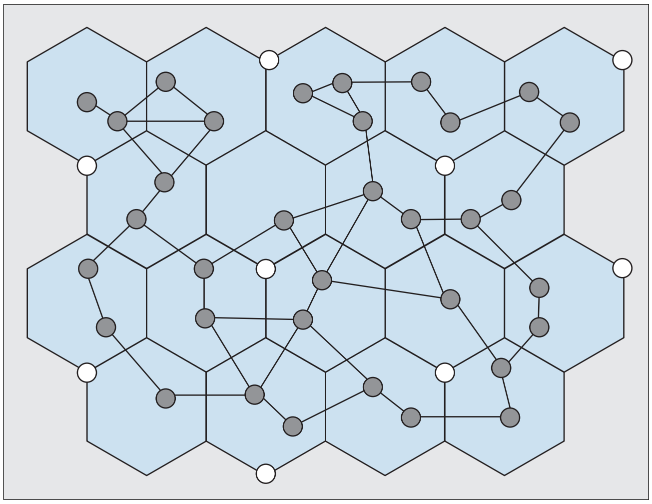

■ Figure 4. *Example of an interference graph.*

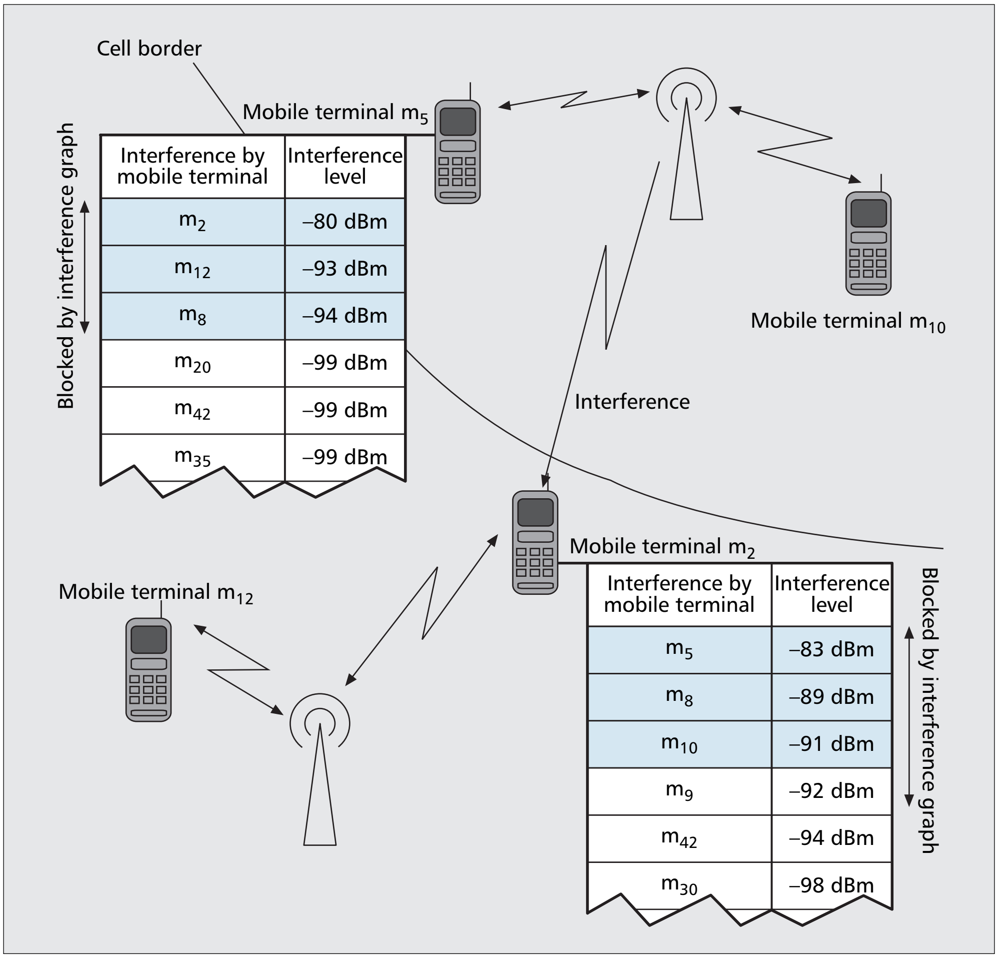

Figure 5. *Creation of an interference graph.*

Performance Metrics

We classify the considered metrics into two groups: area-dependent metrics and area-independent metrics. One of the most important area-independent metrics is the overall spectral efficiency of the system, given in bits/Hz/s. It quantifies how many information bits can be transmitted per second in a certain available bandwidth range and thus, decouples the throughput performance from the available bandwidth. The spectral efficiency is related directly to the aggregate cell sector throughput.

The aggregate sector throughput does not take into account fairness issues. Although it is relatively easy to provide a high data rate to terminals close to the base station, the terminals at the cell borders must not starve. An additional metric that covers this issue and that is commonly used within the 3GPP standardization is the five-percent quantile of the user throughput [14]. In a Monte-Carlo simulation, the five-percent quantile is easy to determine. In our event-driven scenario, we determine the quantile by measuring the individual throughputs of the terminals within a short time period T_S and then calculate the quantile of these short-term averages. An important factor here is the short term period T_S, which must be long enough to even out effects of the MAC procedure, such as segmentation and retransmissions, but short enough to capture coverage holes within the sector area.

In addition to these throughput metrics, the signal-to-interference ratio (SIR) and the signal-to-interference/noise ratio (SINR) are of great interest. Cellular systems with a dense grid of base stations in an urban environment are interference limited. Consequently, noise can be neglected, that is, SIR ≈ SINR. The mean of the experienced SIR is a relatively bad metric because it is dominated by high SIR values close to the base station. If a single scalar value characterizing the SIR conditions is desired, it is better to use the median, which is not susceptible to outliers. Furthermore, the distribution function of the SIR within a cell sector gives important information on the system coverage.

In contrast to area-independent metrics, area-dependent metrics correlate a certain performance value with a particular geographic position. One example is the SIR distribution over the area. Such an area-dependent metric extends the previously discussed fairness metrics in a very illustrative way. For example, it enables the quick assessment and identification of problematic border areas suffering from low SIR conditions. In this article, we define the *observation area* as indicated in Fig. 6 within which we will consider area-dependent metrics.

In a network with *frequency reuse 1*, all cells are allowed to use the full available frequency spectrum, that is, there is full inter-cell interference. With *frequency reuse 3*, only one third of the available frequency spectrum is used in a cell sector. For example, in Fig. 6, cell sectors 0, 3, 6… would use the same frequency band, whereas cell sectors 1, 4, 7… would use disjoint bands. In the following, we consider an uncoordinated reuse 3 system with beamforming antennas as our reference scenario. The considered permutation zone is subdivided into three equally large portions, each of which is assigned to one of the three sectors of a base station. The area-dependent SIR is shown in Fig. 3. The so-configured system achieves a decent

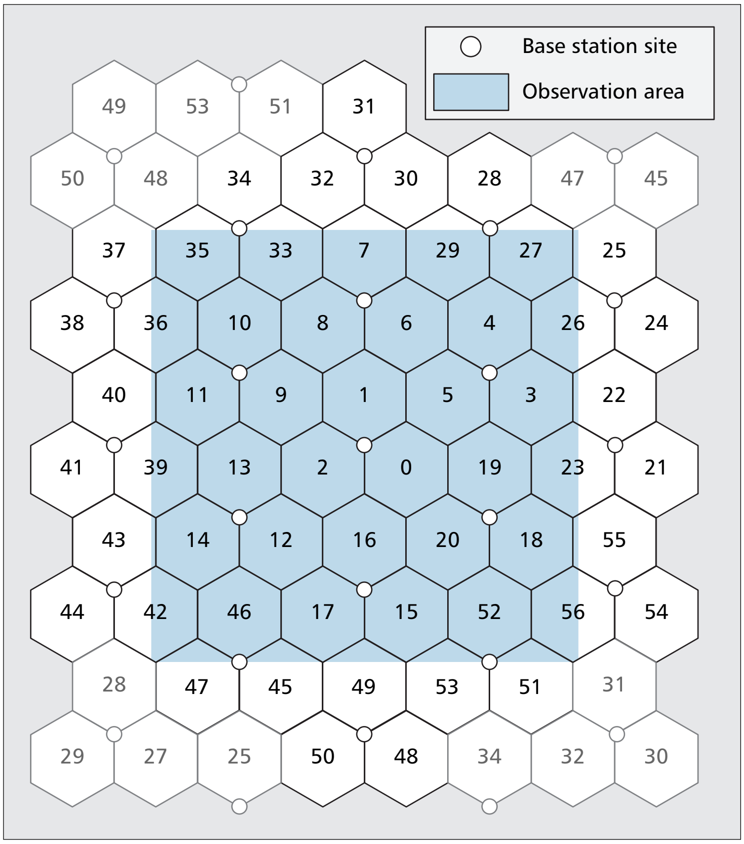

■ Figure 6. *Hexagonal cellular layout with wraparound. The layout consists of 19 base station sites with a total of 57 cell sectors.*

SIR performance throughout the area, with excellent SIR conditions close to the base station. Note that the SIR conditions in a reuse 1 system would be worse, in particular at the sector borders, even at locations close to the base stations [1]. The aggregate sector throughput of the reuse 3 system is about 889 kbit/s, with a five-percent throughput quantile of about 297 kbit/s.

IFCO with Global System Knowledge

A global omniscient entity for interference coordination was applied in ad hoc networks in [11] and in [1] for cellular OFDMA networks. Such an entity is capable of acquiring the global system state instantly and perform scheduling decisions for all base stations on a MAC frame level. Naturally, such an approach cannot be realized, but nevertheless it is interesting to consider for several reasons. First, it can provide an upper bound for the achievable performance of an interference coordinated system. Second, it enables the evaluation of some basic coordination parameters, for example, the coordination diameter (see below), which is an important tool for the development of algorithms that actually can be implemented.

The algorithms in [11] and [1] are based on a conflict graph or an interference graph, respectively. They indicate critical interference relations in between the mobile nodes. In the ad hoc scenario, the conflict graph connects nodes that may not transmit at the same time on the same resource [11]. In the cellular scenario, the interference graph connects mobile terminals that may not be served at the same time on the same resource [1]. Figure 4 shows an example of an interference graph in the cellular scenario. Based on such a graph, the central entity decides on the mobile terminals that are to be served in the next MAC frame on the available resources by the individual base stations. This can involve an additional scheduling algorithm, for example, a scheduler to provide a certain QoS to each connection.

Creation of Interference Graph

The creation of the interference graph can be based on different goals and by using different algorithms. Figure 5 illustrates the algorithm outlined in [1]. This algorithm assures that a given desired minimum SIR (D_S) is achieved for every transmission. For this, an ordered list is generated for every mobile terminal. This list contains the potential interference levels caused by transmissions in other cell sectors to other mobile terminals. Then, the strongest interferers are marked in the list and connected in the interference graph, which avoids simultaneous transmissions on the same resources. The algorithm thus blocks as many of the strongest interferers until the desired SIR D_S is met for every mobile terminal.

The list of potential interferers described above can be generated based on the knowledge of the path loss between the mobile terminals and the base stations. For every mobile terminal, the list must be generated individually. This can be done either by considering all other mobile terminals as interfering terminals (∞-tier coordination), by limiting this list to terminals served by a neighboring base station (1-tier coordination), or by considering only those terminals in neighboring sectors served by the same base station (0-tier coordination). The number of coordinated tiers also is referred to as the coordination diameter.

Simulation Model and Scenario

We consider a hexagonal cell layout comprising 19 tri-sectored base stations at a distance of d_{BS} = 1400 m, which is illustrated in Fig. 6. The scenario is simulated with wraparound according to [15], meaning that the 19 base stations are infinitely repeated in the xy-plane. For example, a cell on the right side of the scenario causes interference to a cell on the left side of the scenario. Therefore, there is no distinct center cell, and all cells are equal. This leads to a completely symmetric and balanced scenario and enables capturing measurement results in all cells, leading to a dramatic reduction in simulation time. Every base station has three 120° cell sectors, where each sector is served by one transceiver. The transceivers are equipped with linear array antennas with gain patterns according to [1]. They can be steered toward each terminal with an accuracy of 1° degree, and all terminals can be tracked ideally. Note that in practice, the steering accuracy is within a few degrees of the actual angle of arrival (AoA) and depends on the scatter environment.

Each sector contains N = 9 mobile terminals moving at a velocity of v = 30 km/h. The underlying mobility model is a random direction model with a mean free path length of 50 m and a maximum turning angle of 25°. As all mobile terminals are bound to their respective cell sectors, they are reflected at the sector borders if they were to leave the sector. A greedy traffic source transmits data towards each terminal, that is, there is always data available to be transmitted for a terminal.

The system model was implemented as a frame-level simulator using the event-driven simulation library IKR SimLib. Path loss was modeled according to [16], terrain category B. Slow fading was considered using a log-normal shadowing model with standard deviation 8 dB. Frame errors were modeled based on block error ratio (BLER)-curves obtained from physi-

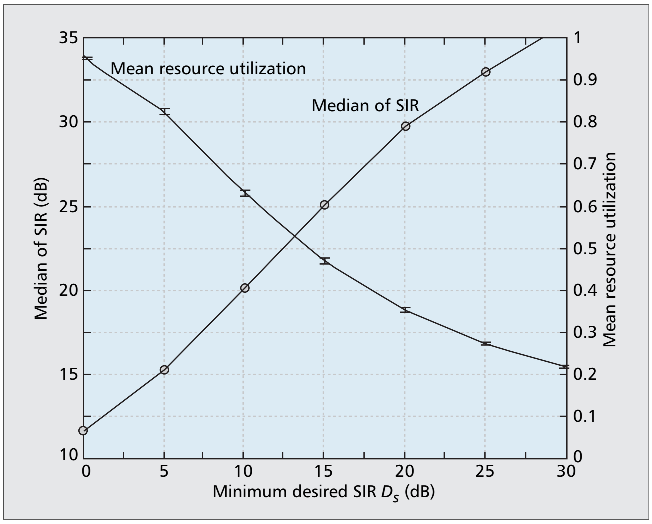

■ Figure 7. *Median of SIR and utilization of resources for a globally interference coordinated system (one-tier coordination).*

cal layer simulations. The simulation model comprised all relevant protocols, such as fragmentation, automatic repeat-request (ARQ) and hybrid (H)ARQ with chase combining. Throughput measurements were done on the Internet Protocol (IP)-layer, capturing all effects of SINR variations and retransmissions.

Performance Evaluation

The presented interference coordination algorithm works with a frequency reuse of 1, which basically allows a resource utilization of 100 percent. However, a higher desired SIR implies a larger number of conflicts in the interference graph, making it more difficult to utilize all available resources. This implies a lower resource utilization, which counteracts the higher SIR, leading to a trade-off between SIR and resource utilization. Figure 7 illustrates this trade-off by plotting the mean resource utilization and the median of the SIR within a cell sector over the desired minimum SIR D_S. Note that the median of the SIR is always larger than D_S because D_S only indicates the minimum required SIR. As a reference, the frequency reuse 3 system has a resource utilization of 33.3 percent with a SIR median of 22.5 dB.

Due to this trade-off, the SIR or the resource utilization are not sufficient metrics, and it becomes necessary to investigate the throughput performance of the system. The aggregate sector throughput is plotted in Fig. 8 (left) depending on the desired SIR D_S and for different configurations. The chart reveals two important facts. First, there is a maximum of the throughput for a particular D_S, which could be expected from the previously discussed trade-off. Second, this maximum is achieved for higher values of D_S as the number of coordinated tiers decreases. The reason is because the algorithm has less control over the cell border areas as the number of coordinated tiers decreases. Consequently, the interference from those base stations, which can be controlled, must be smaller, which can be achieved by a larger value of D_S.

As discussed earlier, the aggregate throughput does not capture fairness concerns in the system, making the throughput quantile an important metric. Figure 8 (right) plots the five-percent quantile over the aggregate value of the throughput. Shown are the best values achieved for different numbers of coordinated tiers. Here, the increase of the aggregate throughput also increases the five-percent quantile. Although the performance of the 1-tier coordinated system is comparable to the ∞-tier coordinated system with respect to the aggregate throughput, it falls short with respect to the throughput quantile.

Interference Coordination with Local System Knowledge

As mentioned before, the coordination by a global omniscient entity is not practical. It is desirable to perform interference coordination within each base station, based on local system knowledge. In fact, the already mentioned 0-tier coordination can be done based on local system knowledge within one base station and therefore, is not only of theoretic interest. However, it was shown in the previous section that the performance of a full ∞-tier coordination cannot be reached. This is mainly due to the algorithm having no control over the SIR in the border areas of the cells. This problem can be alleviated by means of FFR, as was done for example in [8].

Several variations of such a scheme are possible. In [8], the reuse 1 and reuse 3 areas are on disjoint frequency bands, whereas [17] and [18] use the full set of available resources in the reuse 1 areas and one-third of the same resources in the reuse 3 areas. Variations are also possible with respect to the transmit power level in each of the areas. In [17], the reuse 1 areas are covered with a reduced power level, whereas in [18], the transmit power of interfering base stations is reduced.

The assignment of mobile terminals to reuse 1 or reuse 3 areas can be done based on the distance of a mobile terminal from its serving base station or on the present SINR situation. The SINR-based assignment can be done by measurements in the mobile terminal. These can be based on the measurement of pilots from the serving and the interfering base stations or on measurements of recently received data frames. The measurements must be fed back to the base station, which is also required for other purposes, such as burst profile selection.

Performance Evaluation

In this article, we use the full set of resources for the reuse 1 areas and one-third of the same resources for the reuse 3 areas. The power is not controlled as part of the interference coordination but in the course of the burst-profile management, based on the present SINR conditions of a mobile terminal. Mobile terminals are assigned to the reuse 1 or reuse 3 areas, based on the exponential average of their previously experienced SINR values. For this purpose, two SINR thresholds are introduced. If the SINR exceeds the upper threshold, the corresponding terminal is assigned to reuse 1. If it falls below the lower threshold, it is assigned to reuse 3. If the SINR is in between, no changes are made.

Moreover, we apply a coordination in between sectors of the same base station, very much as it was proposed in [8]. In our example, the algorithm for coordination was the same as in the globally coordinated system with 0-tier coordination. Such a coordination can be performed with only local information in a base station and therefore, is well realizable.

Figure 9 (left) shows the SIR distribution within the area for such an FFR system. Compared to the pure reuse 3 system of Fig. 3, the average SIR conditions are worse but acceptable even at the cell border. The SIR conditions must be worse because the reuse 3 and reuse 1 areas share the same resources in our FFR system. If they used disjoint frequency ranges, the SIR conditions would be better than in the reuse 3 case, but the resource utilization would be worse.

The SIR conditions that are worse compared to the reference reuse 3 system are compensated by a better resource utilization. Figure 9 (right) shows the resource utilization within the observation area. It shows how the mobile terminals are assigned to the reuse 1 area more often as they become closer

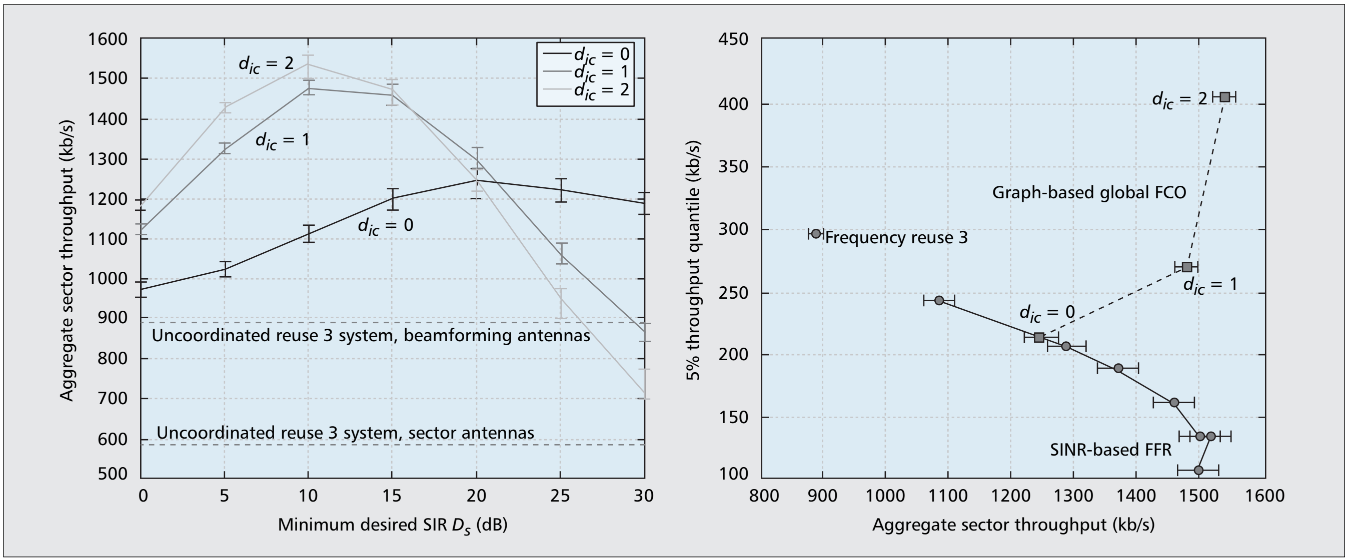

■ Figure 8. *Aggregate sector throughput over desired SIR for global coordination (left) and 5 percent throughput quantile over aggregate sector throughput for different configurations (right).* $T_S = 4$ *s.*

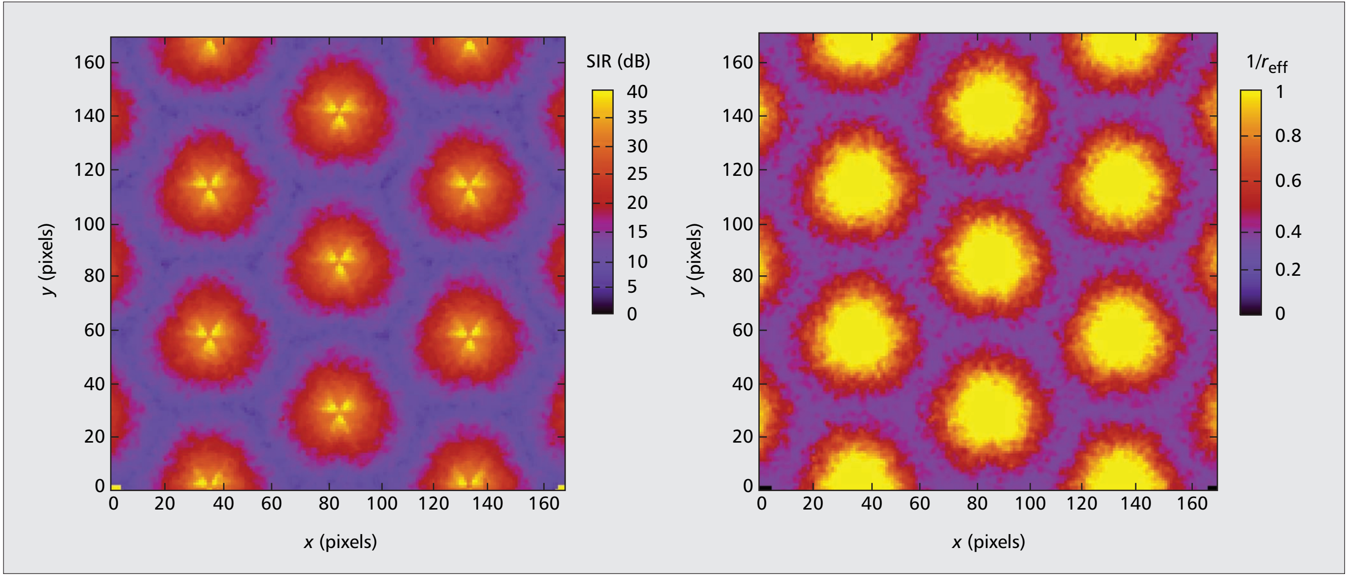

■ Figure 9. *SIR conditions (left) and resource utilization (right) with fractional frequency reuse (FFR).*

to the base station. In fact, the reuse 1 area covers a relatively large portion of the coverage area, leading to a good average resource utilization of 66 percent.

Last but not least, the FFR system shows very good performance with respect to the aggregate sector throughput, as shown in Fig. 8 (right). The plotted curve shows the performance of the FFR system for the different choices of the SINR thresholds. It can be seen that the aggregate and the cell edge performance can be traded off against each other by the SINR threshold setting. Thereby, the FFR system can match the globally coordinated system with respect to the aggregate throughput. However, the five-percent throughput quantile is much lower. Note that the quantile must be even lower than in the pure reuse 3 system due to the sharing of reuse 1 and reuse 3 resources in the FFR system, which leads to the SIR degradation at the cell border that already was observed. Finally, we note that FFR is a very attractive approach to mitigate inter-cell interference but lacks mechanisms to improve the fairness with respect to the cell border areas.

Conclusion

In this article, we provided an introduction to interference coordination in emerging OFDMA systems. After an overview of relevant literature, we discussed several important metrics to evaluate the performance of interference coordination algorithms. Subsequently, to obtain an upper bound estimate of the system performance, we introduced an interference coordination algorithm that is based on global system knowledge. We then discussed possibilities to achieve interference coordination that could be well implemented with only local-system knowledge within each base station.

Although it is always easy to boost the aggregate cell throughput at the expense of fairness, the actual challenge is to maximize the aggregate throughput and the cell border throughput at the same time. We showed that a global interference coordination algorithm can achieve both quite well. On the other hand, the local algorithm can almost match the global scheme with respect to the total aggregate cell throughput but falls short with respect to the throughput in the outer areas of the cell sectors.

Acknowledgments

This research was performed in cooperation with Bell Labs Germany, Stuttgart. The author would like to thank Bozo Cesar, Stephan Saur, and Andreas Weber for their valuable input and discussions.

References

[1] M. C. Necker, "Towards Frequency Reuse 1 Cellular FDM/TDM Systems," *Proc. 9th ACM/IEEE Int'l. Symp. Modeling, Analysis, Simulation Wireless and Mobile Sys.*, Torremolinos, Malaga, Spain, Oct. 2006, pp. 338–46.

[2] WiMAX Forum, "WiMAX Forum Mobile System Profile," rel. 1.0, May 2007.

[3] 3GPP TSG RAN WG1#44 R1-060586, "Downlink and Uplink Intercell Interference Co-ordination/Avoidance — Impact on the Specifications," *Tech. Rep.*, Ericsson, NTT DoCoMo, Denver, CO, Feb. 2006.

[4] D. C. Cox and D. O. Reudink, "Dynamic Channel Assignment in High-Capacity Mobile Communications Systems," *Bell Sys. Tech. J.*, vol. 50, no. 6, July–Aug. 1971, pp. 1833–57.

[5] J. Engel and M. Peritsky, "Statistically-Optimum Dynamic Server Assignment in Systems with Interfering Servers," *IEEE Trans. Commun.*, vol. 21, no. 11, Nov. 1973, pp. 1287–93.

[6] S. Halpern, "Reuse Partitioning in Cellular Systems," *Proc. IEEE VTC'83*, vol. 33, Toronto, Canada, May 1983, pp. 322–27.

[7] T. Bonald, S. Borst, and A. Proutière, "Inter-Cell Scheduling in Wireless Data Networks," *Proc. Euro. Wireless'05*, Nicosia, Cyprus, Apr. 2005, pp. 566–72.

[8] M. Sternad *et al.*, "Attaining Both Coverage and High Spectral Efficiency with Adaptive OFDM Downlinks," *Proc. IEEE VTC'03-Fall*, vol. 4, Orlando, FL, Oct. 2003, pp. 2486–90.

[9] R. Vilzmann, C. Bettstetter, and C. Hartmann, "On the Impact of Beamforming on Interference in Wireless Mesh Networks," *Proc. IEEE Wksp. Wireless Mesh Networks*, Santa Clara, CA, Sept. 2005.

[10] R. Ramaswami and K. Parhi, "Distributed Scheduling of Broadcasts in a Radio Network," *Proc. IEEE INFOCOM'89*, vol. 2, Ottawa, Canada, Apr. 1989, pp. 497–504.

[11] K. Jain *et al.*, "Impact of Interference on Multi-Hop Wireless Network Performance," *Proc. 9th Annual Int'l. Conf. Mobile Comp. Networking*, San Diego, CA, 2003, pp. 66–80.

[12] Y. Wu *et al.*, "Bounding the Power Rate Function of Wireless Ad Hoc Networks," *Proc. IEEE INFOCOM'05*, vol. 1, Miami, FL, Mar. 2005, pp. 584–95.

[13] P. Stuedi and G. Alonso, "Computing Throughput Capacity for Realistic Wireless Multihop Networks," *Proc. 9th ACM/IEEE Int'l. Symp. Modeling, Analysis, Simulation Wireless and Mobile Sys.*, Torremolinos, Malaga, Spain, Oct. 2006, pp. 191–98.

[14] 3GPP TS 25.814, "Physical Layer Aspects for Evolved Universal Terrestrial Radio Access (UTRA)," *Rel. 7*, June 2006.

[15] IEEE 802.20 Working Group on Mobile Broadband Wireless Access, "IEEE 802.20 Evaluation Criteria — Version 17R1," Sept. 2005.

[16] V. Erceg *et al.*, "An Empirically Based Path Loss Model for Wireless Channels in Suburban Environments," *IEEE JSAC*, vol. 17, no. 7, July 1999, pp. 1205–11.

[17] 3GPP TSG RAN WG1#42 R1-050841, "Further Analysis of Soft Frequency Reuse Scheme," *Tech. Rep.*, Huawei, London, U.K., Sept. 2005.

[18] 3GPP TSG RAN WG1#42 R1-050764, "Inter-Cell Interference Handling for E-UTRA," *Tech. Rep.*, Ericsson, London, U.K., Sept. 2005.

Biography

MARC C. NECKER (marc.necker@ikr.uni-stuttgart.de) received an M.S. degree from the Georgia Institute of Technology in 2001 and a Diploma degree in electrical engineering and information technology from the University of Stuttgart, Germany, in 2002. He is currently the head of the wireless research group at the Institute of Communication Networks and Computer Engineering at the University of Stuttgart. His research interests are in end-to-end and system performance of wireless networks, modeling of wireless networks, and architecture of future radio access networks. He is participating in several privately and publicly funded projects in the area of mobile networks, and he is a TPC member of several international workshops and conferences.