Local Negotiation in Cellular Networks: From Theory to Practice

Raz Lin and **Daphna Dor-Shifer** and **Sarit Kraus**
Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel 52900
{linraz, shiferd, sarit}@cs.biu.ac.il

David Sarne
Division of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138
sarned@eecs.harvard.edu

Abstract

This paper describes a novel negotiation protocol for cellular networks, which intelligently improves the performance of the network. Our proposed reactive mechanism enables the dynamic adaptation of the base stations to continuous changes in service demands, thereby improving the overall network performance. This mechanism is important when a frequent global optimization is infeasible or substantially costly. The proposed local negotiation mechanism is incorporated into a simulated network based on cutting-edge industry technologies. Experimental results suggest a rapid adjustment to changes in bandwidth demand and overall improvement in the number of served users over time. Although we tested our algorithm based on the service level, which is measured as the number of covered handsets, our algorithm supports negotiation for any set of parameters, aiming to optimize network's performance according to any measure of performance specified by the service provider.

Introduction

A cellular network is a radio network made up of base stations. The base stations in the cellular network are used for radio communication with mobile dynamic agents. The cellular network aims to provide service to the mobile clients, each according to their service level agreement (SLA). This is done by optimizing some or all of the network's parameters, such as: coverage, cell power level, pilot fraction, capacity, antenna's tilt of the base station and others (Hampel *et al.* 2003). The coverage of each base station (BS) varies according to the terrain, the location of the station, landmarks, interference by other stations and the antenna parameters (such as tilt). Also, a power change by each BS can cause a change in the users served by the BSs, including itself.

The infrastructure of existing cellular networks is mostly a static one. RF engineers are involved in the testing and planning of the (possibly) optimal deployment of the network using commercial optimizing tools, after which the network is deployed. However, as the number of users and bandwidth demands rise (e.g. due to rich media deployed services), the complexity of optimizing the network performance increases as well. This becomes particularly relevant, when migrating to the third-generation (3G) and fourth-generation (4G) networks (Rappaport 2001). While in the first-generation (1G) and second-generation (2G) networks, the requested bandwidth could be associated with a voice service, i.e. fixed requested bandwidth per call, the introduction of data services in 3G and 4G networks significantly increases the variance in bandwidth demands. Also, the uncertainty regarding the number of users that needs to be served in the network at a given time and the services each such user requires (for example, downloading a movie, or simply making a phone call) make it difficult for current cellular networks to react in real time to the users' online demands and optimize the network. Finally, the accumulated experience of recent years underscores the effect of evolving events, such as crowded events (e.g. the Super Bowl) or an unexpected environmental catastrophe (e.g. Hurricane Katrina) on network performance. These have short-term as well as long-term implications for network overload, and thus require real-time responses.

Recent advances in cellular technologies suggest new opportunities in terms of real-networks' time control and dynamic configuration of base stations. One dominating technology is the "smart antenna" enabling dynamic changes to the parameters of the BS. However, there is a need to provide methods that will apply the new capabilities to meet the emerging needs. Today, once the network is deployed, most of its configuration changes are done at predefined times and by using global optimization tools, such as Schema's UMTS OptiPlanner (Schema 2005), rather than being dynamically triggered by real-time demand changes coming from the users' side. This kind of global optimization is resource consuming and costly (for example, Schema's optimization requires a few hours to optimize a medium-sized cellular network). Thus, a central optimization cannot be executed frequently, and new innovative techniques are required to allow a better utilization of the network and load-balancing adjustments of the traffic.

We propose a novel approach to be incorporated into the infrastructure of cellular networks to improve the performance of the network, bypassing the need for frequent, and usually infeasible, global optimizations. Our approach is a reactive approach, in that it enables, via negotiation, changes in the base station's parameters (e.g. the change in the pilot power of the base stations, which is its total transmission power). The negotiated changes are generated and evaluated according to their predicted effect on the network's performance as measured, for example, by load balancing the network's traffic or the global coverage of the network. Though the evaluation process of the negotiated changes cannot precisely predict the effect on global network performance (due to the localization of the calculation and parallel negotiations that take place) overall, our proposed mechanism improves network's performance over time. Furthermore, in case of conflicting changes, the mechanism immediately recovers by readjusting the relevant BSs' configuration. The distributed nature of the method implies several important advantages such as minimized communication cost and the ability to quickly combine partial information to form a good global assessment (Shen, Zhang, & Lesser 2004; Yadgar, Kraus, & Ortiz 2003). Although the implementation described in this paper focuses on negotiating a single parameter, the negotiation protocol itself can be easily extended to allow the negotiation over any set of parameters.

A second contribution of our research is the introduction of one of the first integrated simulation environments for cellular networks with an agent-oriented paradigm. Most simulation environments for cellular networks have thus far focused on technical implementation and modeling of the cellular network, such as: propagation model, traffic distribution, path loss and others (Rappaport 2001). We have succeeded in incorporating agent-oriented capabilities, i.e. transforming the base stations to reactive, autonomous, entities with AI capabilities, which can serve as a test-bed for numerous aspects of artificial intelligence and agent-based mechanisms in cellular networks, far beyond the negotiation protocol.

The distributed negotiation mechanism, as well as the system we describe, were developed as part of the RAN Optimization group in the REMON consortium, targeted at the development of pre-competitive generic technologies for the fourth-generation (4G) Mobile Cellular Systems.

Key Challenges

The transformation of the static base stations into reactive AI entities presented us with several challenges regarding the negotiation protocol that we decided to implement. These can be categorized as challenges that influence the theoretical-based aspects of our algorithm and challenges that influence the empirical implementation and testing of our algorithm in a realistic cellular network. The challenges associated with the first group included:

- **Selecting the negotiation protocol:** A cellular network is usually a very large distributed network. The antenna resource, which is utilized for communication purposes, should not be excessively used for purposes other than communications between mobile users. Thus, a negotiation between all the base stations in the network, resulting in sending numerous negotiation messages using the antenna resource, seems to be too costly, whereas distributed local negotiation appears to provide a better solution (Du *et al.* 2003; Du, Bigham, & Cuthbert 2003; Kraus 2001; Mailler, Lesser, & Horling 2003; Shen, Zhang, & Lesser 2004; Xuan, Lesser, & Zilberstein 2001). However, using a local negotiation mechanism might cause only minor local changes which have no impact on the global network, or it might cause conflicting local improvements, which overall, worsen the network's performance. The designed protocol should overcome these potential problems.

  Note that although we propose a distributed negotiation mechanism, it does not necessarily imply that the changes in the network will be done in a distributed manner. For example, since the base stations in today's cellular networks are connected to a centralized system (e.g. Mobile Telephone Switching Office), which is responsible for coordinating the base stations and providing handoff operations, each base station can be modeled as an entity in that system. Our proposed distributed negotiation mechanism enables the negotiation between those entities, while the actual change in the network is managed by the centralized system.
- **Neighborhood Formation:** For a distributed local negotiation, the communication is to be made with predefined (or dynamically defined) negotiation partners ("neighbors"). A decision on the method for tagging other base stations as neighbors should be carefully made since it has considerable impact on the mechanism's performance. The following aspects of neighborhood formation should be addressed: (a) Should the neighborhood be based on the locations of the base stations (for example, geographical distances or geographical location)? (b) Should it rely also on the level of interference? (c) Is there room to incorporate the level of influence between the base stations, caused by changing the power of other antennas as a parameter in this process?

In the second category, we find challenges that impact the method that allows the empirical evaluation of our proposed mechanism:

- **Integration with cellular network simulation:** As we wanted to compare the usefulness of our approach to real cellular networks, we needed to incorporate the agent-based mechanism in an existing cellular network simulation. The incorporation of two such complex systems, in order to generate a single integrated system, presented us with many technical, as well as design challenges. Ideally, we would have tested our mechanism on an existing cellular simulation tool. However, while surveying existing simulation tools, such as Andrew's Odyssey (Andrew) and Actix's CellOpt ACP (Actix), we learned that none of them support adjustable autonomy at the base station level. We therefore we had to rely on an "off-the-shelf" state-of-the-art optimization tool - Schema's UMTS OptiPlanner(Schema 2005) - to model the network, and develop a system that would interface with this tool to model the base stations as self-contained entities.

In the remainder of this paper, we will present our approach given the above challenges, both in terms of the theory and implementation. In the next section, we will review related work in the field of distributed local negotiation, and continue with the presentation of our negotiation protocol. Then, we will describe the experimental settings and the results. Finally, we will conclude and propose future work in this field.

Related Work

As stated above, a negotiation of all the agents with each other is highly costly. This is also supported by (Xuan, Lesser, & Zilberstein 2001), who argue that while the communication is crucial for the coordination of the different agents, it is unrealistic for the agents to reach perfect communication.

In the context of cooperative negotiation, (Shen, Zhang, & Lesser 2004) studied the relationship between the degree of local cooperation, the characteristics of the environment and the global utility achieved by all agents in the negotiation. Their statistical analysis shows that mechanisms for local negotiations, that will allow the optimization of the system dynamically, can be designed. (Mailler, Lesser, & Horling 2003) present a negotiation model for the task of resource allocation in soft real-time environments, in which the agents are both autonomous and cooperative. They show that the cooperative nature of the agents makes it possible to maximize the social utilities of the agents. These two papers motivated us to design the distributed local negotiation mechanism, to enable the efficient utilization of the network, and also make it possible to reach near-optimal global optimization, using local changes.

(Du *et al.* 2003) present a local negotiation approach triggered by the congestion level in the network. The trigger is made by the base station itself whenever it observes that its utilization exceeds a given threshold. Their results show that local negotiation is effective for the network and yields a performance very close to that obtained by global optimization techniques. However, as opposed to (Du *et al.* 2003), we aim to allow the negotiation to be made at any given time and not depend on a single agent's view of the network load. While we are simulating real cellular networks, their simulations had some permissive assumptions concerning the cellular networks which they simulated. For example, they assume that the interference comes only from other traffic units in the same cell, whereas in real cellular networks, as in our simulations, this is not the case.

(Du, Bigham, & Cuthbert 2003) present a utility-based approach for geographic load balancing in mobile cellular networks. The cooperation is encapsulated in the utility function rather than in exchanging negotiation messages. The utility function proposed is employed on a traffic unit, that is, a user that generates traffic to the network, and is composed of the total traffic load at each base station and the distance of the traffic unit from the given base station. The utility determines whether a traffic unit is served at a given base station or at another base station. We, on the other hand, try to optimize the network by making it reactive to real-time events and not to each single traffic unit in the network. As such, we allow the network to serve any traffic unit at any given time, while preserving the load-balancing of the network.

In the next section, we will describe our negotiation protocol and the mechanism of the offer evaluation.

The Negotiation Protocol

We propose a bilateral negotiation scheme between the base stations in which each base station is capable of negotiating with its neighbor stations. The negotiation is done over a change in the configuration parameter/s of the base stations (for example, the pilot power of both negotiators). As we stated above, this can be easily extended to support negotiation over a set or a subset of the network's parameters. Formally, let V denote the set of possible values for a given parameter of the base stations, $v_i, v_j \in V$, O denotes the set of possible offers such that $o(i,j) = ((v_i, v'_i), (v_j, v'_j)) \in O$ is an instance of an offer made by BS named i to BS named j, indicating the change in the parameter value for BS i and BS j, respectively, wherein v_i and v_j are the current parameter values and v'_i and v'_j are the new parameter values. The negotiation itself is local - both for the agents doing the negotiation and in the evaluation of the offer. Notice that in our proposed distributed mechanism, the overall bandwidth, i.e. cost, required for communication is negligible in comparison to the amount of overall bandwidth existing today in 4G, and even 3G networks. This is due to the fact that only local negotiation messages between two base stations are transmitting at each iteration.

We continue in the next section with a description of the distributed local negotiation mechanism incorporated into the cellular network.

Distributed Local Negotiation Model

In order to allow local negotiation between the base stations, we needed to define how a local evaluation of the negotiation can be done so that it best reflects the potential global change in the network's performance. To this end, we defined a locality for each base station and the local evaluation of the offers. The locality of each agent consists of two levels of neighborhood. The first level, denoted L_1, includes the immediate neighbors of the agent. We denote by $L_1(i)$ the set of neighbors in the L_1 level for agent i. The second level, denoted L_2, includes the second level of neighbors, that is, the immediate neighbors for all the L_1 neighbors. Formally, $L_2(i) = \{L_1(j) | j \in L_1(i)\}$. An example is described in Figure 1. Each agent can send an offer to its neighboring agents. Each neighbor agent evaluates the sent offer based on its neighbors and then returns the evaluation result to the sender. The sender then evaluates the offer based on its neighbors and the returned results. The evaluation is based on the utility function, which is described below.

The utility function evaluates an offer locally. Since we propose a general model, the utility function can change, according to the network's performance measure that we want to optimize. We present here an example of such a local utility function which, as stated, aims to optimize the network's performance based on the number of covered handsets. This utility function calculates the difference between the served mobiles before the change in the parameter value and after the change in its value. Formally,

Figure 1: Levels of neighborhood: Agent A's point of view

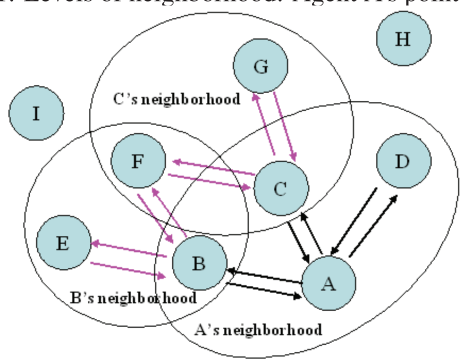

let $servedMobiles_j(v_j, v_i)$ denote the number of mobiles served by a given BS j with a parameter value of v_j and a neighbor BS i with a parameter value of v_i. To enable each BS calculate the number of mobiles it can serve, the BS has to take into consideration the pilot power of its neighbors as well as the distribution of the mobile users in its locality. To this end, at each iteration, each BS transmits its pilot power, if it was changed from the previous iteration, and distribution of users in its area to all of its neighbors. Note that evaluating the effect on the network caused by a change in a network's parameter is not an easy task. This is due to the fact that even a change in a single network's parameter can influence the connectivity of the mobile users in the network and the interference between the base stations. Fortunately, Schema's application, supplies an efficient evaluation tool for this purpose, which returns a quick evaluation to the number of total served mobiles by a given base station. This tool is used in our simulations to calculate the absolute and relative differences in the service level parameter described above. Finally, let $u : O \rightarrow \mathbb{N}$ be the utility of an offer to be implemented, which is defined as:

$$\begin{aligned} u_j(o) &= u_j(\,(v_i, v_i'), (v_j, v_j')\,) \qquad (1) \\ &= servedMobiles_j(v_j', v_i') - \\ &\quad servedMobiles_j(v_j, v_i) \end{aligned}$$

We denote by $U_j(o)$ the sum of utility values of all the L_1 neighbors of j, including j itself, from an offer o. Let BS i be the proposer of an offer o. In order to evaluate the offer, i calculates the utility values of all its L_1 neighbors, while each L_1 neighbor calculates the utility values of all its L_1 neighbors that are distinct from the L_1 neighbor of the proposer i. Formally, let $val : O \rightarrow \mathbb{R}$ be the value of an offer $o \in O$, calculated by the proposer i of the offer, then:

$$\begin{aligned} val_i(o) &= u_i(o) + \sum_{j \in L_1(i)} U_j(o) \qquad (2) \\ &= u_i(o) + \\ &\quad \sum_{j \in L_1(i)} [(\sum_{k \in L_1(j), k \notin L_1(i)} u_k(o)) + u_j(o)] \end{aligned}$$

In the next subsection we address the neighborhood formation problem.

Neighborhood Formation

In the context of deciding on the negotiation's partners for each agent, (Baert & Semé 2003) illustrated the importance of correctly identifying those partners in cellular networks. Too many neighbors might cause a large communication overhead, while too few neighbors might allow for only small local changes. Also, the neighbors can be constructed by taking into account several network parameters, and thus, choosing the best parameters can greatly influence the performance of the algorithm and that of the network.

In our simulations, we tested the negotiation protocol based on several neighborhood definitions, and the results obtained indicate that the user-served-threshold based method generates the best result for the network's performance, measured by the number of covered handsets. In this method, for each BS a change of a specified parameter in a given range is made and all other BSs are checked to see how they are affected. All of the BSs which had a change of at least $\pm T$ are considered to be neighbors of the given BS. This method best reflects the relations between the BSs, yet fine tuning of the threshold is needed. We tested the change in the pilot power parameter and set the range to 25-35dBm[1], T was set to 5 and we looked at the effect of the number-of-mobiles-served parameter.

The Negotiation Sequence

The negotiation itself can be triggered by different events. Examples of such events include global events (e.g. time interrupt) or local events (e.g. a base station observes that the number of served mobiles exceeds a certain threshold). The specific trigger to be used is principally external to our proposed mechanism. Specifically, in our simulations we used predefined time-unit intervals as the trigger for initiating a local negotiation. Each time unit of the negotiation consists of several synchronized serialized phases, in which a given set of actions can be made. The phases by their order in a given iteration are listed below:

1. *Proposal Generation*: In this phase, one base station[2] s can generate offers and send them to its neighboring stations. In our simulations a random base station was selected in each iteration to generate the proposals. The base stations generated three proposals in each iteration. Each such proposal consisted of a random change in the pilot power parameter. Note that this setting was used only to prove the applicability of our method. Obviously, this can be significantly improved by implementing more efficient heuristics in the base station to generate offers tailored for the specific environment it is operating in.
2. *Returned Offers*: In this phase, each offer sent to BS r is returned to the sender s with the evaluation of its value to r, based on the utility value of r and its L_1 neighbors.

[1] dBm represents a measured power level in *decibels* relative to 1 milliwatt. It is used to express power.

[2] This can be easily extended to allow for any number of base stations to generate offers in each iteration.

3. *Evaluation Phase*: This phase consists of an evaluation process, carried out by the sender s, of the returned offers in order to select the best offer. As we have mentioned, this phase involves using Schema's evaluation tool to evaluate the number of served mobiles in the network.
4. *Commitment Phase*: In this phase, the sender s selects an offer for commitment, based on a comparison between the best offer it initiated and the best offer it has received.

The base stations can exchange the following messages, based on the negotiation phase:

- *Offer Messages*: An offer including the change in the parameter value for the proposer and the change in the value for the designated base station.
- *Evaluation Messages*: In this message, an evaluation of a proposed offer is sent back to the proposer.
- *Commit Messages*: After a base station chooses the offer that is best for it, it sends a commit message to the base station that is involved in that offer.

The negotiation protocol is based on hand-shakes, that is, an offer obtains commitment at the final phase of each negotiation's iteration if both agents decide to commit the same offer.

Schema's UMTS OptiPlanner (Schema 2005) was used as a simulation tool, both prior to the simulations and during them. Prior to the simulations it was used in order to create simulated cellular networks, which are replications of real cellular networks. Throughout the negotiation, it was used for evaluation purposes. Using its evaluation tool, we could evaluate the effect of the different offers on the cellular network, and using our utility function, decide which offer will potentially produce the best improvement to the network.

We note that the local negotiation mechanism we have presented above simplifies the process of making changes in the network. Even with our proposed simple mechanism we managed, in a distributed manner, to reach a solution, not too inferior to the optimal solution, which is significantly timely and resource consuming if done, otherwise, centrally.

Experiments

In order to convey the advantages of our proposed mechanism to the cellular industry, so that it will embrace this approach, we demonstrated the mechanism's capabilities in realistic simulations. To this end, we applied a leading industrial optimization tool for real cellular settings (Schema 2005). Schema's UMTS OptiPlanner, a tool with field-proven experience, is a centralized automatic base-station planning and optimization solution that produces optimal base-station parameter configurations, based on user-defined goals for quality, capacity, coverage and budgetary constraints. OptiPlanner optimizes a wide range of key network configuration parameters, such as: antenna location, antenna height, type, tilt, azimuth and power settings. Furthermore, as we have described above, OptiPlanner has a fast evaluation tool that takes all the aspects of real cellular networks into account when measuring the performance of a given network.

Based on this tool, we have developed an agent-based simulation environment, which integrates both the cellular network and the negotiation architecture. Our simulations enable the checking of various methods for the dynamic adaptation of cellular systems' parameters. Each time unit of the simulation simulates a complete negotiation iteration, in which entities that simulate the base stations negotiate together over a local change of the specified parameter of the pilot power. This change will enable global improvement in the network's performance.

Experiment Settings

We tested our proposed model on a circular-shaped network model scenario and a snake-shaped network model scenario. Figures 2(a) and 2(b) demonstrate the clutter (land cover affecting propagation loss for each base station) and terrain of the circular and snake shaped models, respectively. The circular-shaped model scenario consists of 30 base stations and an average network radius of ten kilometers, while the snake-shaped model scenario consists of 27 base stations and an average network radius of seven kilometers. The cellular network itself is a model of a real network deployed in the suburbs of a large European city. The use of these two distinct models allows us to examine the efficacy of our proposed mechanism on different networks with varying degrees of impact between the base stations. For example, in the circular-shaped model scenario, the effect on the base stations is different: while the stations at the center of the circle are affected similarly from all sides, those on the exterior show diminished results.

The configuration of each scenario is built using Schema's UMTS OptiPlanner tool (Schema 2005). This basic network configuration includes the propagation model, clutter type and topographic map. Other parameters are also defined, using the simulation tool and real data, such as: traffic load, pilot power and the antenna parameters of the base stations. Given this data, other parameters for the antenna are defined or calculated, such as: the 3D pattern of the antenna, its type, gain, frequency and tilt. The average number of users per sector and total path loss are calculated based on all these parameters. These configurations and parameters genuinely reflect cellular networks, which are deployed throughout the country. During the negotiation process itself, the base stations negotiate over the different values for the pilot power that were set during the configuration process.

The simulations were designed to analyze the efficacy of the distributed local negotiation, when dynamic changes in the density of the network and in the distribution of the mobile users occur. Each simulation involved more than 300 iterations and a total of 150 minutes.

During the simulation, changes in the distribution of mobile users were initiated every 15 iterations, and the density of the mobile users changed every 50 to 80 iterations. The proposals themselves consisted of changes in the pilot power of the base stations. The evaluation of each proposal also involved using Schema's tool as an evaluation tool, which takes some of the network parameters into consideration.

The following subsection presents the results of those experiments.

Figure 2: Clutter and terrain of the (a) circular-shaped model, and the (b) snake-shaped model scenario.

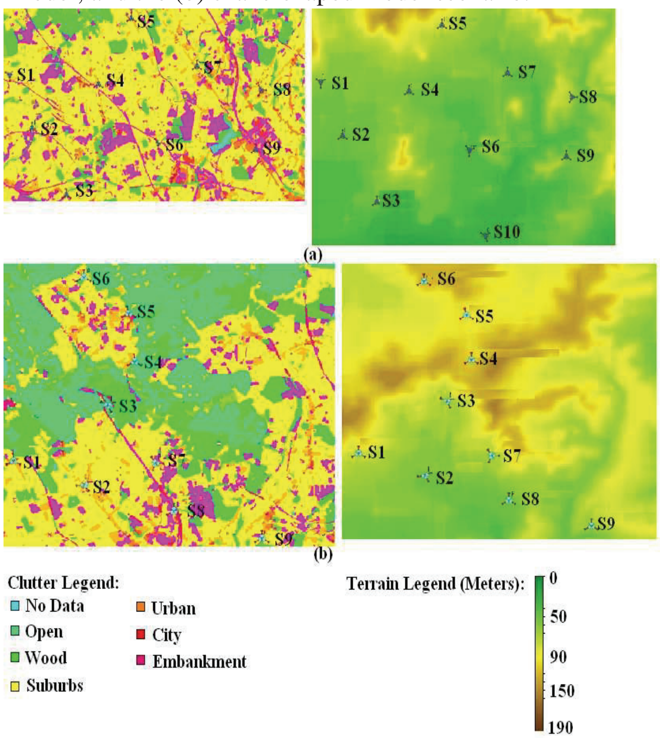

Experiment Results

As stated above, the purpose of the experiments was to model a dynamic environment of cellular users in an urban area, and test the response of our local negotiation mechanism to the dynamic changes. Figures 3 and 4 display the average percentage of mobile users served by our algorithm, in comparison to the performance obtained in the same network after it was initially optimized by the central optimizer for the snake-shaped and the circular-shaped models, respectively.

In a cellular network in the suburbs of a large city of the type considered here, the typical density of mobile users is 2, and consists of approximately 1,200 to 1,500 mobile users. Schema's optimization tool was used to perform an optimization based on this density. Given this optimization, the system can serve, on average, 90% of the users. However, even though the average is 90%, there are usually small variations in the users' numbers and location. So, for a given scenario, the current static cellular system may perform below average. These variations are modeled via what we refer to as a seed change. Each seed yields a different specific setting of density 2. Our negotiating agents adjust to these small changes by negotiating over readjustments of the power pilot of the different BSs they represent. In Figures 3 and 4, this scenario is modeled, for example, in iterations 1-50. The seed change is marked by red squares. The blue line, which specifies the percentage of users served by the cellular system using our negotiating system, is above the pink one, which represents the percentage of users served by the current static system.

Figure 3: 27 base stations, snake-shaped model Results.

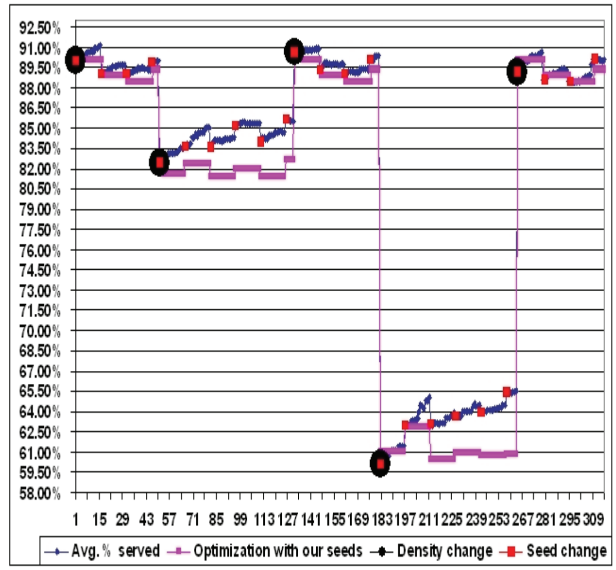

Even though the cellular infrastructure in the suburbs of the big city was built to serve users of density 2, there are situations in which the density is changed, e.g., due to unusual events. We first considered a scenario in which the density is doubled and is 4, with approximately 3000 users. In such a situation, it is clear that the infrastructure will not be able to serve 90% of the users in this situation, as captured in iterations 51-130 in Figures 3 and 4. If no change is made to the power pilot (as is the situation in current static cellular systems) the percentage of users that are served is about 80%, as indicated by the pink line in the figures. However, if our negotiating system is applied, the percentage increases in less than 20 iterations, which is equivalent to less than 15 minutes in the simulation, to 84%. Also, in this scenario the specific number of users and their location changes over time, and this is modeled again by the seed change. Even though the percentage increase is only of about 4%, the total increase in the number of mobiles served at the end of the simulation is about 44,000 for the 30-circle scenario and 16,000 for the 27-snake scenario. The results show that our negotiating system is capable of adapting quickly and successfully to these changes. Even when the resources become scarce, the dynamic nature of our model allows rapid adaptation and enables fast recovery in the percentage of served mobiles.

Once the density is back to normal, that is, density 2 (see iteration 130 in the figures), the current static system immediately returns to the previous average of 90%. Our negotiating system adapts quickly to this change too and adapts the power pilot to the original situation - all this without any outside intervention. However, as opposed to the static system, our negotiation system will keep on adapting to the changes in the number of users and their changing locations. These results indicate that even when the state of resources in the network becomes abundant once again, our algorithm works

Figure 4: 30 base stations, circular-shaped model results.

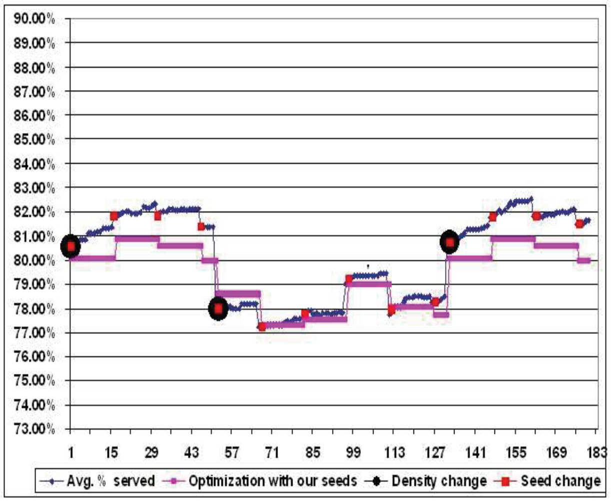

well enough and again serves almost the same percentage of served mobiles, as if there had been no change in resources in the first place.

In addition, we also considered a situation in which the density is changed to 10, with approximately 5,000 users. This can be viewed in iterations 181-260 in Figure 3. The results show that the static system can serve on average about 62% of the mobile users, whereas our negotiating system can reach the level of 70%. Again, when the density returns to 2, our negotiating system quickly readjusts to this change and returns to serve 90% of the mobile users quickly.

Conclusions

This paper demonstrates the promise embodied in integrating a distributed local negotiation mechanism for cellular network simulations, which can affect the real-time adaptation of deployed cellular networks. We have shown the applicability of our proposed method in two distinct scenarios, which bolsters our confidence with regard to its efficacy regarding general scenarios. Obviously, the introduction of efficient heuristics to produce the negotiated offers themselves will significantly improve the performance of the mechanism. This is also true for the incorporation of changes applied to a set of parameters as part of each local negotiation, thus significantly increasing the magnitude of improvement in next generation networks.

Our innovative approach in the integration of simulation environments for cellular networks with an agent-oriented paradigm will allow future test-bedding for other purposes, far beyond the negotiation protocol.

Future work on this field includes the introduction of clusters and intra-cluster negotiation, as well as the investigation of the nature of the committed offers in order to gain a better understanding of the dynamics of the network. As noted, our algorithm can be extended to multi-attribute negotiations, and heuristics can be employed to take more than one parameter into account. Our negotiation protocol initiated a change in the pilot power between two base stations. When investigating the results, it seemed that better performance was achieved when the sum of differences between the pilot power before and after the change, for the negotiation agents, was positive. This issue requires further in-depth investigation and could improve our results.

Acknowledgments

This work was supported in part by the Israeli Ministry of Industry and Trade under project REMON and ISF grant number 8008. We thank Moti Ritz, Shalom Mu'alem and Haim Pistner for their help in understanding the infrastructure of the cellular network and to Morris Ben-Zaken and Arie Reichman for presenting us with the challenges of the REMON project. We thank Michael Livschitz and Tomer Armarnik for their help in understanding and using Schema's optimization tool. We also thank Saar Rosenberg, who contributed to the development of the simulations environment.

References

Actix. http://www.actix.com/products/acp.htm.

Andrew. http://www.andrew.com/products/measurement_sys /odyssey/.

Baert, A.-E., and Semé, D. 2003. An optimal one-to-all fault-tolerant broadcasting in voronoï cellular networks. In *PDCS 2003*, 69–74.

Du, L.; Bigham, J.; Cuthbert, L.; Nahi, P.; and Parini, C. 2003. Intelligent cellular network load balancing using a cooperative negotiation approach. In *WCNC2003*, 1675–1679.

Du, L.; Bigham, J.; and Cuthbert, L. 2003. Utility based distributed geographic load balancing in mobile cellular networks. In *Fourth International Conference on Mobile Communication Technologies, IEE 3G'2003*, 58–62.

Hampel, G.; Clarkson, K. L.; Hobby, J. D.; and Polakos, P. A. 2003. Tradeoff between coverage and capacity in dynamic optimization of 3g cellular networks. In *IEEE Vehicular Technology Conference*.

Kraus, S. 2001. *Strategic Negotiation in Multiagent Environments*. Cambridge MA, USA: MIT Press.

Mailler, R.; Lesser, V.; and Horling, B. 2003. Cooperative negotiation for soft real-time distributed resource allocation. In *AAMAS 2003*, 576–583. Melbourne: ACM Press.

Rappaport, T. 2001. *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ, USA: Prentice Hall.

Schema. 2005. http://www.schema.com/products_umts.html.

Shen, J.; Zhang, X.; and Lesser, V. 2004. Degree of local cooperation and its implication on global utility. In *AAMAS 2004*, 546–553.

Xuan, P.; Lesser, V.; and Zilberstein, S. 2001. Communication decisions in multi-agent cooperation: Model and experiments. *Fifth International Conference on Autonomous Agents* 616–623.

Yadgar, O.; Kraus, S.; and Ortiz, C. L. J. 2003. Hierarchical information combination in large-scale multiagent resource management. In *Communication in Multiagent Systems*, 129–145.