Towards the Fourth Generation of Cellular Networks: Improving Performance Using Distributed Negotiation

Raz Lin and Daphna Dor-Shifer and
Saar Rosenberg and Sarit Kraus
Computer Science Department
Bar-Ilan University
Ramat-Gan, Israel
{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 programmatic approach to efficiently distribute resources in a dynamic cellular network, using local negotiations. Our proposed mechanism is reactive and facilitates parallel self-adaptation efforts, leading to dynamics that improve overall network performance. The local nature of the negotiations being performed as part of the adaptation process enables frequent changes in the network's parameters with a negligible coordination overhead. The results of our experiments suggest rapid adjustment to changes and overall improvement over time in the number of users served by the network. We evaluate our algorithm based on the service level index, measured by the number of covered handsets. Nevertheless, the proposed algorithm supports any set of parameters and any combination of performance measures supplied by service providers.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed applications*; I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Distributed Artificial Intelligence*

General Terms

Algorithms, Experimentation

Keywords

cellular networks, 3G, 4G, distributed negotiation

1. INTRODUCTION

Recent developments in cellular communication technology (e.g. the introduction of third generation networks) has resulted in a high paced emergence of various new portable devices and cellular services. These have been accompanied by a phenomenal increase in end users' demands for increased connectivity and bandwidth, primarily for supporting data services. The nature of these services (e.g. navigation data, video) requires new standards of response times and transmission volumes in communication between base stations and users' portable devices. Furthermore, many of the services are currently provided through existing cellular network infrastructures, which are upgraded in order to comply with newer stiffer expectations (e.g., support of heterogeneous technologies, demand for high bandwidths and interactive network applications).

Expected trends of 4G networks include the convergence of services through IP Multimedia Subsystems (IMS), where a mobile subscriber is continuously online and reachable via all traditional services (e.g., telephony, email, instant messaging and web) and is also able to operate them. In addition, the 4G networks provide a variety of services, each of a different type, that users require (e.g. downloading a movie or simply making a phone call). These advances in the domain of large cellular networks present challenges, as well as opportunities: from the design and deployment of new antenna types up to network maintenance, resource utilization and ongoing adaptation and adjustment [9].

Meeting the goal of maintaining increasing demands for faster communication rates and Quality of Service (QoS), can not be obtained solely by improving the infrastructure of the network since a certain level of cooperative relations between the different elements in the network is also required. For example, a municipal wireless mesh network can be used to relieve some of the load from the overall cellular network, and, similarly, residential networks can be used to handle some of the traffic. This might even require the cooperation of potential foreign service providers[1] (and even different types of cellular networks) to assure the high level of QoS the users demand in such a competitive industry. Also, in order to enable a media independent handover and/or other types of resource sharing, a certain degree of abstracted cooperation is required.

In [11] we introduced intelligent software agents, designed for operating on behalf of base stations in a cellular network thus enabling the base stations to act as autonomous, yet

[1] A cooperation between different foreign service providers is not unrealistic. As resources become scarce, and less antennas shall be permitted, different service providers could benefit from interacting with one another.

MSWiM'06, October 2–6, 2006, Torremolinos, Malaga, Spain.

cooperative, entities. Each of these entities has the capabilities to negotiate with other entities in order to evaluate local changes in the network's parameters, which could possibly improve the overall performance of the network. This also allows dynamic changes in the base stations, as is expected in the 4G networks.

Our preliminary results in [11] have shown that our mechanism enables, via local negotiations, improvement in the performance of large cellular networks with dynamic changes. In this paper we extend these results and offer some fine-tuning to our suggested algorithm, alongside the introduction of a new heuristic for generating proposals for negotiation. We continue and analyze the results and provide insights from our experience in modeling realistic cellular networks for test-bedding purposes.

In addition to the encouraging results our software agents have achieved, we introduce one of the first integrated simulation environments for cellular networks with an agent-oriented paradigm. While most simulations use models, such as: the propagation model, traffic distribution and path loss [14], we have built a realistic reconstruction of a deployed cellular network, while demonstrating the benefits of reactive intelligent agents in settings as close to reality as possible. Thus, our simulation environment can serve as a test-bed for numerous aspects of artificial intelligence and agent-based mechanisms in cellular networks, far beyond our proposed negotiation protocol.

The distributed nature of the AI agents implies several important advantages, such as: minimized communication cost and the ability to quickly combine partial information to form a good global assessment [17, 20]. Although the implementation described in this paper focuses on negotiating a single parameter in a 3G cellular network, the negotiation protocol itself can be easily extended to allow the negotiation over any set of parameters in any network.

The distributed negotiation mechanism, as well as the system we describe, have been developed as part of the RAN Optimization Group of the REMON Consortium, targeted at developing pre-competitive generic technologies for 4G Mobile Cellular Systems.

This paper is organized as follows. In Section 2 we review related work in the field of distributed local negotiation. Section 3 describes our negotiation protocol and the mechanism of the offer evaluation. In Section 4 we describe the simulation environment, our experimental settings and results. In Section 5 we propose a heuristic for the offer creation, and show experimental results for its efficacy. Finally, we conclude and propose future work in this field.

2. RELATED WORK

In recent years greater focus has been placed on developing dynamic mechanisms that will enable re-adjustment of cellular networks to their environment in real-time. For example, in Bell Labs [5] researchers experimented with dynamic control algorithms that can automatically adjust the network to changes in both traffic and network conditions, and adapt autonomously when new cells are added to the system. They argue that real-time measurements are a fundamental ingredient in the development of dynamic control mechanisms. Essentially, this is due to the ability to reveal inefficiencies in the network by gathering information on the actual traffic characteristics and users' behavior.

In addition to real-time data gathering, a proposed algorithm can also benefit from a distributed mechanism. This mechanism can provide a simple way of making local decisions that can affect the entire network. However, the distributed mechanism should be devised carefully. Negotiation of all the agents (base stations) with each other is highly costly in terms of the utilization of the network. This is also supported by Xuan et al. [19], 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 et al. [17] 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 allow optimization of the system dynamically, can be designed. Mailler et al. [12] 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 use of cooperation enables social welfare maximization. Though they did not focus on cellular networks, these two papers motivated us to design the distributed local negotiation mechanism that enables the efficient utilization of the network by constantly maintaining near-optimal global optimization, using local changes.

Du et al. [7] present a local negotiation approach triggered by the congestion level of the network. Here, each base station initiates a negotiation with its neighbors upon observing its utilization exceeding a pre-defined 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. [7], our algorithm is service-based rather than network state based. We aim to allow the negotiation to be made at any given time without depending on a single agent's view of the network load. Furthermore, while we simulated 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, interference is caused by neighboring base stations.

Du, Bigham and Cuthbert [6] 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 react 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.

3. THE NEGOTIATION PROTOCOL

We recall from [11] that the bilateral negotiation scheme between agents enables reaching an agreement, which will improve the overall network performance. Each agent, representing a base-station, is capable of negotiating with a list of other agents in the network, termed "negotiation partners" or "neighbors" (neighbors are further discussed in Section 3.2). The negotiation is done over a change in the configuration parameters of the other base stations (for example, the pilot power of both negotiators). As we stated above, this method can easily be extended to support negotiation over a set or a subset of the network's parameters, and negotiation between other types of physical entities. 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 (base station) 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 conducting the negotiation and in the evaluation of the offer. Note that our protocol enables the distribution both in the negotiation level (the algorithm) and in the physical level (the distributed algorithm can be incorporated into each of the network's nodes, rather than placing all the agents in a central location). Thus, our algorithm can be used in a spectrum of applications. For example, when a central physical implementation is selected (such as the Network Control Center (NCC) or the Mobile Telephone Switching Office (MTSO), which is responsible for coordination of the base stations and providing handoff operations to the mobiles while they move), the network's statistics can be directly accessed. This allows each agent to return its computation results to the manager entity which decides when and how to implement the changes. However, the disadvantage of a central entity is the lack of concurrency in calculations amongst the agents. Conversely, the distribution mechanism enables concurrent calculations at a messaging cost, which is negligible as compared to the overall bandwidth required for communication and data services in today's 3G networks. This is due to the fact that the local negotiation messages between two base stations only include a minimal payload (e.g., numeric data on the pilot power and calculations of the changes in traffic as a result of altering it).

The following subsections detail technical implementation issues, which enabled us to evaluate our proposed mechanism in real cellular network settings.

3.1 Distributed Local Negotiation Model

In order to allow local negotiation between the base stations, we also needed to define how a local evaluation of the negotiation can be accomplished so that it best reflects the potential global change in the network's performance. To this end, in [11] 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 neighborhoods, based on either definition of neighbors, as described in the next subsection. 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, $L_2(i) = \{L_1(j) | j \in L_1(i)\}$. This is illustrated in Figure 1. Each agent can send an offer to its neighboring agents. Each neighboring 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. This process can be repeated continuously for a predefined number of iterations.

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

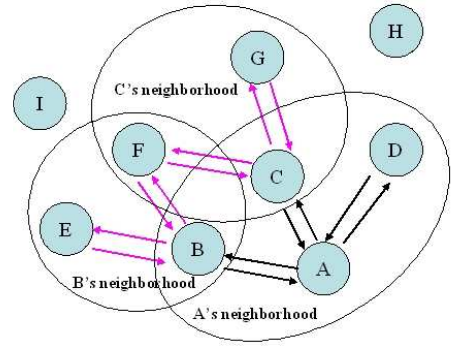

The utility function evaluates an offer locally. Since we propose a general model, the utility function can be changed according to the network's performance measure that we want to optimize. We present an example of such a local utility function when the optimized parameter is pilot power (in our experiments we also optimized the azimuth parameter), and this function, as stated, aims to optimize the network's performance based on the number of covered handsets. This utility function calculates the difference between the mobiles served before the change in the parameter value and after the change in its value. Formally, 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 to 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 the distribution of users in its area to all of its neighbors. Note that evaluating the effect of 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. For example, if a base station increases its pilot power, then it may now have more interference with its neighboring base stations, which may cause the disconnection of some of the users in its area of service. Fortunately, our network simulation application supplies an efficient evaluation tool for this purpose, which returns a quick evaluation of the total number of mobiles served 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:

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

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 neighbors 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:

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

In the next subsection we describe the different possibilities of forming neighbors in the context of cellular networks.

3.2 Neighborhood Formation

In the context of deciding on negotiation partners for each agent, Baert et al. [3] illustrated the importance of correctly identifying partners in cellular networks. Too many neighbors might cause a large communication overhead, while too few neighbors might result in small local changes. Also, the neighbors can be constructed by taking into account several network parameters. Thus, choosing the best parameters can greatly influence the performance of the algorithm and of the network. In the cellular network's infrastructure, the base stations generally maintain a neighbors' list, which is based on the hand-over state of the mobiles. While using these neighbors' lists seems as the trivial option, it appears that it would indeed be futile. The main cause for this is the purpose for which the neighbors' list was created, as compared to the purpose of the neighbors of each BS for negotiation purposes. For example, when constructing the neighbors' lists some neighbors may be overlooked because the location of mobiles was not optimal while testing for the neighborhood, and therefore did not reflect the actual connections between the BSs. To this end, we considered several neighborhood definitions, such as: (i) full connected network, (ii) geographical based neighborhood, (iii) hand-over state neighborhood (that is, for each mobile agent a list of best intercepted BSs, according to the signal strength, is maintained), and (iv) threshold based users served, in which 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 underwent 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. In our simulations, when we tested the change in the pilot power parameter and set the range to 25-35dBm[2], T was set to 5 and we looked at the effect of the number-of-mobiles-served parameter. The electrical azimuth parameter was changed in the ranges of $\pm 25^\circ$ from the current location. For this parameter the threshold was set to 2. It is interesting to note, that when we considered electrical azimuth, some of the BSs had only themselves as a neighbor. This means that a change in their own azimuth only affects the users in their area of service, and do not affect any of the users of other base stations. Therefore, a proposal with a change in their own azimuth value can be sufficiently evaluated by the BS itself.

In the cellular networks' infrastructure, there is also the issue of a symmetry of the neighbors. We assert that BS_i is a *symmetric neighbor* of BS_j if BS_i is a neighbor of BS_j and BS_j is also a neighbor of BS_i. Although it may seem intuitively true that neighborhoods are always symmetric, our experience in the context of cellular networks, proved this assumption wrong. Therefore relying on inherent symmetry rather than the definition given above results with additional communication overhead and complication of the algorithm's calculations, without generating better solutions.

In our simulations, we tested the negotiation protocol based on all the different neighborhood definitions that were given above. 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 paper we focus on the experiments in which our negotiation mechanism was incorporated with neighborhood relations that were defined using the threshold based method, since this method seemed to be the most beneficial one. The results of these experiments are presented in Section 4.

3.3 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 agent (base station)[3] s can generate offers and send them to its neighboring agents. In our simulations a random base station was selected in each iteration to generate the proposals. The base station generated three proposals in each iteration. Each proposal consisted of a random change in the selected parameter (e.g. pilot power). 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, as described in Section 5.

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

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

2. *Returned Offers*: In this phase, each offer r sent to BS is returned to sender s with the evaluation of its value for r, based on the utility value of r and its first-degree neighbors.

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 mentioned, this phase involves using the integrated 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 negotiation protocol is based on hand-shakes, that is, an offer obtains a commitment at the final phase of each negotiation's iteration only if both agents decide to commit to the same offer. Schema's UMTS OptiPlanner [16] was used as a simulation tool, both prior to and during the simulations. Prior to the simulations it was used in order to create simulated cellular networks, which are replications of real cellular networks. Throughout the negotiations, 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 would potentially produce the best improvement for the network.

We note that the local negotiation mechanism 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 near optimal solution, which is significantly timely and resource consuming if, otherwise, done centrally.

4. EXPERIMENTS

We ran several sets of experiments in order to test our proposed distributed mechanism. Each set varied in the status of the network when the experiments began (a semi-optimized or optimized network state), in the parameters optimized during the negotiation and in the environment (the number of antennas and their deployment). In the next subsections we describe our simulation environment, the experimental settings and the experiments' results.

4.1 The Simulator

In order to verify the applicability of our proposed mechanism, we demonstrated the mechanism's capabilities in realistic simulations. Ideally, we would have tested our mechanism on an existing cellular simulation tool. However, while surveying existing simulation tools, such as Andrew's Odyssey [2] and Actix's CellOpt ACP [1], we learned that none of them support adjustable autonomy at the base station level. We therefore had to rely on an "off-the-shelf" state-of-the-art optimization tool - Schema's UMTS OptiPlanner [16] - to model the network, and develop a system that would interface with this tool to model the base stations as self-contained entities. Schema's UMTS OptiPlanner [16], a tool with field-proven experience, is a centralized automatic base-station planning and optimization solution. The tool 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 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.

As mentioned earlier, the entities negotiating in a distributed manner to improve the network's overall performance can either be placed on a centralized location (e.g. on computers placed in the MTSO, which is responsible for coordination of the base stations and providing handoff operations to the mobiles while they move) or in distributed locations (on the base stations themselves). Either way, the distributed nature of our proposed mechanism remains. Placing it in a centralized location enables minimization of the antenna resources needed for communication between the autonomous agents, since it can now be done using the MTSO [9].

In order to show the efficacy of our algorithm, we needed to achieve a dynamic nature in the static simulator. Thus, in the simulation we initiated three types of changes between consecutive time frames in which our algorithm operates, as follows:

- *Small alternations in traffic:* the users' distribution remains the same and a similar number of users (not necessarily the same ones) requires service. For example, this setting can act as a reflection of users in a shopping mall when during certain hours a constant distribution of users remains, yet their identity changes, as shoppers come and go quite frequently. In our simulations this case was simulated by a change in the seed parameter, which represents the user distribution.

- *Large alternations in traffic:* in this type there are changes in the number of users requiring services and in their distribution. This situation can reflect the workplace at the beginning (or at the end) of a workday, when many workers arrive or leave their workplace at once, causing the bandwidth requirements to change. In our simulations this situation was simulated by a change in both the seed and the density[4] parameters.

- *Constant traffic:* in this type no changes are made in the overall number of users' requesting service nor in their distribution. This simulates a situation which is almost static. This can be seen, for example, in a workplace during the working hours, when almost all of the mobile users remain in a relatively constant location. In our simulations this case was simulated by not varying the seed nor the density between consecutive time frames.

[4]The density is used to simulate different loads of the network (e.g., density 2 means that we double the amount of traffic of the initial environment, and the new traffic is distributed with similar proportions as in the initial environment).

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

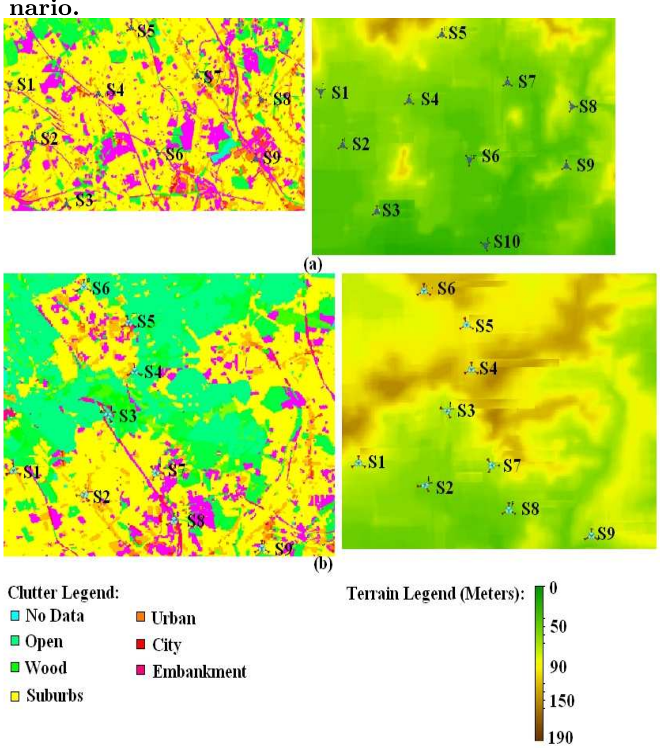

4.2 Experiment Settings

The cellular network we use is a model of a real existing network deployed in the suburbs of a large European city. In these experiments we looked at two network scenarios, each with different characteristics : a circular-shaped network scenario and a snake-shaped network scenario. Figures 2(a) and 2(b) demonstrate the clutter[5] and terrain of the circular and snake shaped scenarios, respectively. The circular-shaped scenario consists of 30 base-stations with an average network radius of ten kilometers, while the snake-shaped scenario consists of 27 base stations, with an average network radius of seven kilometers. 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 27 snake-shaped scenario there is less impact between base stations located at the extremes, while 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 [16]. This basic network configuration includes the following:

- *Propagation model*, based on (a) the Deygout model [15] augmented by improvements similar to the ones added to the Walfisch-Ikegami model [8, 10] at the COST 231.
- *Clutter types* which include a representation of area obstacles (possible types in the software include: water, open, light woods, woods, tall woods, wooded suburbs, low suburbs, suburbs, city and more)
- *Topographic map*, a genuine terrain's heights map.

Other parameters of the network are also defined using the simulation tool and real data, such as: traffic load[6] (in the simulations we refer to changes in the traffic load as *seed changes*), pilot power and the antenna parameters of the base stations (such as the height it was installed and its pattern).

Given this data, other parameters of the antenna are defined accordingly, and their resulting impact on the network expected performance is calculated, such as: the 3D pattern of the antenna, its gain frequency, its horizontal and vertical width, and its electrical tilting ability.

Based on all these parameters, the average number of users per sector and the total path loss can be calculated. These configurations and parameters genuinely reflect cellular networks, which are deployed throughout the country. Before the negotiation process begins, a parameter is chosen to be optimized, and is given a range of possible values. During the negotiation process itself, the base stations negotiate over the selected parameter in the given range.

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. We allowed each iteration to last 30 seconds. In each simulation we measured the number of mobiles being served, whereas each mobile is considered served if the base stations can provide services for it, based on its service level agreement (SLA).

As explained earlier, changes were initiated as to the distribution and density of users throughout the experiments. 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 generated by the base stations to their negotiation partners consisted of proposed changes to the pilot power or to the azimuth of the base stations (depending on the parameter selected to be optimized). The evaluation of each proposal is based on the previous time unit and involves using Schema's tool as an evaluating tool, which takes some of the network parameters into consideration. This is done in a continuous manner to simulate the expected real behavior of the network.

The following subsection presents the results of those experiments.

4.3 Experiment Results

Recall that the purpose of the experiments was to test if our mechanism leads to an improvement in the service level of cellular users, in a dynamic environment in an urban area. In order to show the generality of our suggested mechanism we tested two different parameters of base stations in the network: pilot power and azimuth. The azimuth, though

[5] Land cover affecting propagation loss for each base station.

[6] The amount of traffic carried out by the sector during the period that is being simulated, and is either based on data from the switch or on marketing estimates. The traffic units are measured in Erlang.

not ordinarily changed, can serve as a good example to our tool's flexibility in incorporating different aspects of a complicated system, which may have been otherwise neglected. It is also a parameter that is enabled for electrical tampering in MIMO [4]. These two parameters potentially offer a significant impact on the network in terms of changes in the coverage or in the service areas, resulting in improvement of the overall network performance.

We note that the results of the dynamic algorithm we use have an accumulated improvement in service level over the results of the global optimization, though there are some local isolated points in time in which there is a decrease in the service level, as compared to the global optimization. We attribute this to the *ripple effect*, a common phenomenon in cellular networks, in which the change in one part of the network affects other parts of the network that were not changed [13]. Due to this effect, the dynamic locally-positive changes which are carried out in one part of the network, may occasionally cause interference in other parts of the network and eventually cause an overall degradation in network performance. However, this effect is minimized in our simulations, due to our neighborhood selection process, which takes into account the effect on other BSs that might be influenced by a change in network parameters (see Section 3.2).

The subsequent subsections describe the results of our simulations on the different network parameters. First we present the results when negotiating over a change in the pilot power, and then we continue with the results of negotiating over a change in the azimuth of the base stations.

4.3.1 Negotiating over Pilot Power

We use the term *pilot power* to refer to the CPICH common pilot power. As explained in [18], pilot power optimization in UMTS networks is a crucial engineering issue which has received increasing interest during the last few years. Optimizing the P-CPICH power levels facilitates coverage control and significantly improves radio resource utilization in UMTS networks. Combined with optimization of radio base stations' antenna tilts, this becomes a powerful tool that allows to significantly reduce the total interference level in the network and to improve the network capacity in a very efficient way. Another characteristics of the changes in the pilot power is that it affects mainly the down-link channel, that is, the link between the base station and the mobiles.

During negotiations, each BS can propose a change in the pilot power in the range of 25dB to 35dBs, which are realistic parameters for the network used. 25-35dB represents a large enough scope to describe most of the valid ranges for these particular parameters from a physical point of view and from experience in the field.

In a cellular network in the suburbs of a large city of the type considered here, the typical density of mobile users is 2, consisting of approximately 1,200 to 1,500 mobile users. Schema's optimization tool was used to perform an optimization based on this density. This was the baseline of all our experiments, density-wise. 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. Therefore, 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.

We ran experiments with different baselines in respect to the network parameters. In [11] we described the first experiment, in which we started out with a fully optimized network. To obtain this baseline we used Schema's OptiPlanner tool, which was run for several hours, with all parameters enabled for use in the optimization process. Results and details on this experiment can be found in [11]. In the following experiment we used a semi-optimized network, i.e. a network that is not fully optimized, but underwent a partial optimization process. This was done in order to ensure the performance of our algorithm also in a non-optimized, or random, situation - as may be the case when the global optimization is already outdated when the dynamic algorithm is launched. The density used for the optimization in both cases was 2.

In the experiment in which we used a less optimized baseline scenario, all BS's have a pilot power of 33. This point was reached after a semi-optimization process using Schema's OptiPlanner, in which only the pilot power parameter was given as a possible parameter for optimization. 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 semi-optimized by the central optimizer for the snake-shaped and the circular-shaped models, respectively.

Figure 3: 27 base stations, snake-shaped model results for partial optimization.

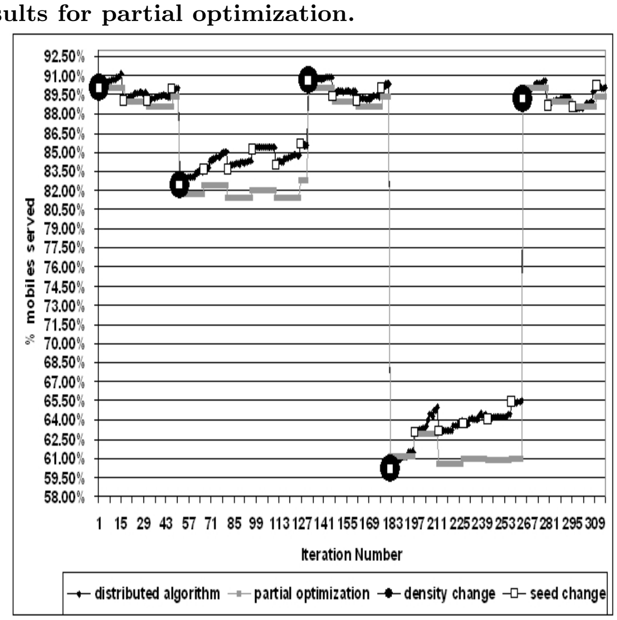

Given this optimization, the system can serve, on average, for example in the 27-snake scenario, 90% of the users. As can be seen, once again our negotiating agents adjust to the small seed changes by negotiating over readjustments of the power pilot of the different BSs they represent. In Figures 3 and 4, this scenario is modeled, for instance, in iterations 1-50.

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

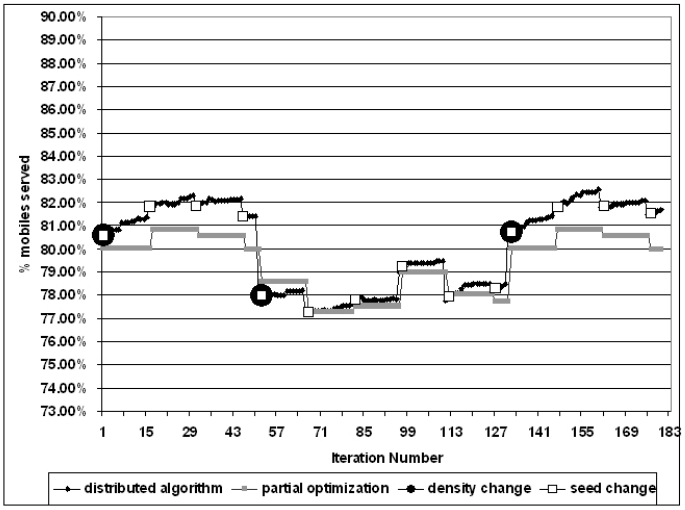

Figure 5: 27 base stations, snake-shaped model results for full optimization using density 10. Optimizing azimuth.

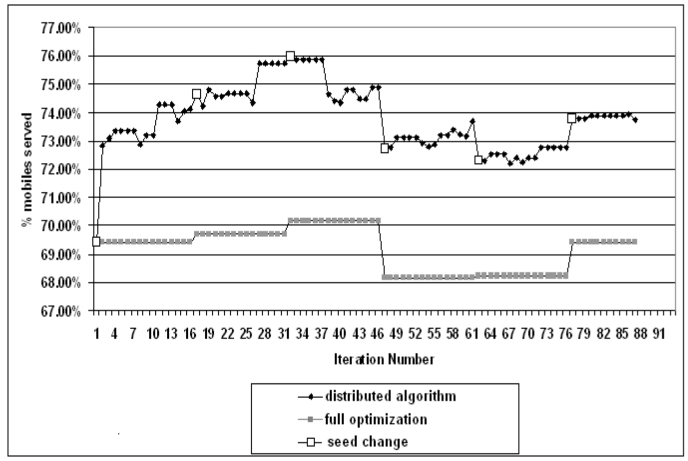

As in the previous experiment, we now raised the density to 4, with approximately 3000 users. If no change is made to the pilot power (as is the situation in current static cellular systems) the percentage of users that are served is about 80%. However, if our negotiating system is applied, the percentage increases in less than 20 iterations to 84%, which is equivalent to less than 15 minutes in the simulation. This is an improvement in the percentage of users being served, and it reflects the changes that occurred in traffic within this time period.

We continued and lowered the density back to normal (density 2, see iteration 130 in Figures 3 and 4), and the outcome demonstrates that the current static system immediately returns to the previous average. Our negotiating system quickly adapts to this change and will keep adapting to the changes in the number of users and their changing locations. Afterwards, the density was changed to 10, with approximately 5,000 users. This can be viewed in iterations 181-260 in Figures 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 almost instantly.

The results show that our negotiating system is capable of successfully adapting to the ongoing changes in traffic and recovering from momentary drops. This enables constant improvement in network resource utilization.

In the next subsection we present the experiment results of applying our algorithm using a change in the azimuth rather than the pilot power parameter.

4.3.2 Results of Negotiation on Azimuth

To demonstrate the generality of our algorithm and its ability to improve the network's service level when negotiating on different parameters, we also used the electrical azimuth as a negotiated parameter (as we have mentioned above, this is in addition to the experiments in which the negotiation was done over changes in the pilot power). The electrical azimuth used enables the base station's antenna to rotate left or right according to a given value. The base stations negotiated over a change of $\pm 25°$ from the current value.

Electronically shifting the base stations to the different azimuths locally agreed on by their agents, can lead to an improvement in the overall service level. A change in the azimuth can lead to an improvement in the percentage of mobiles being served since, although there are three antennas on a pole, they are not situated back to back. Also, the interference levels of two base stations serving the same area are not necessarily strong enough to prevent an improvement in service level. A proposal generated by each base station consists of a change in only one antenna on the pole.

Our baseline was a network which has undergone full optimization using Schema's OptiPlanner on density 2. In Figure 5 a fully optimized network immediately received a traffic load of a density of 10. As can be seen from the results there is an added improvement in the service level over the full optimization when the density is 10. The dynamic algorithm suggests an average increase of 4.68% improvement in each iteration. However, when running the algorithm on a network with a traffic load of a density of 2, there was no improvement of the dynamic algorithm over the global optimization. This is expected, because the network resources are still abundant, and the global optimization should be the best setting available for this traffic. When the traffic changes from what the global optimization expects and there is competition on the network resources, there is a great improvement of the dynamic algorithm over the global optimization, as indeed can be seen in Figure 5.

Our result show that our mechanism is not limited to a single parameter, but rather it is general and any parameter can be used for negotiation.

5. HEURISTIC AND DISCUSSION

In order to gain a better understanding of the dynamics of the network we investigated and analyzed the nature of the committed offers in the experiments where the pilot power

Figure 6: Mobile power delta sum for 27 base stations, snake-shaped model, based on full optimization.

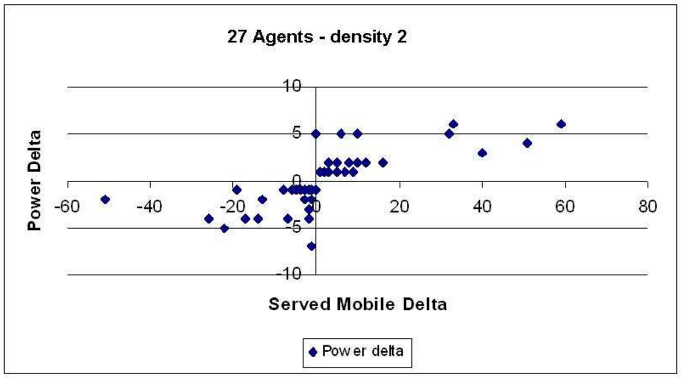

Figure 7: 27 base stations, results of heuristic with a density of 10, scenario based on full optimization.

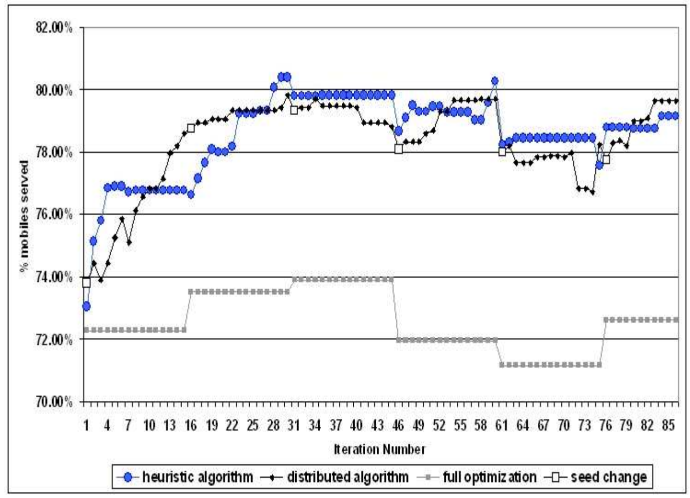

Figure 8: 27 base stations, results of heuristic with a density of 4, scenario based on full optimization.

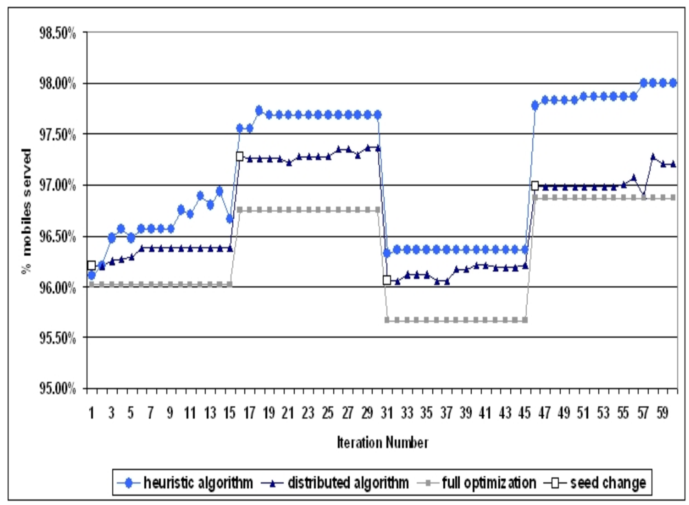

parameter was changed. When investigating the results, we saw that better performance was achieved when the sum of differences between the pilot power before and after the proposed change by the base station was positive. To this end, we introduce the term *power delta sum*, which means the total of the differences in the agent's power changes as a result of a proposal. For example, if the power of BS_i and BS_j is currently 34 and 29 respectively and BS_i proposes an offer which changes the pilot power of BS_i and BS_j to 33 and 32, respectively, then their delta sum will be $(33 - 34) + (32 - 29) = 2$.

The y-axis and x-axis in the graph depicted in Figure 6 represent the power delta sum and the served mobile delta, respectively. The graph reveals the relation between the power-delta sum and the improvement in the amount of served mobiles as a result of a commitment (i.e. an offer having been accepted and its agreement having been realized). As can be seen, while the upper-right quadrant, which represents offers with a positive delta-sum, has a positive change in the number of served mobiles as a result of accepting the offers, the bottom-left quadrant, which represents offers with a negative delta-sum, has a negative impact on the total number of served mobiles in the network.

Using this insight, we implemented a heuristic based on this *power delta sum*. Basically, a base station will propose an offer only if the power delta sum in its offer is positive. However, if no proposals with a positive power delta sum could be generated, a random proposal would be generated and evaluated. Also, since there are also situations in which a negative power delta sum generates an improvement in the performance of the network, only 90% of the proposals are generated using the heuristic, and the other 10% are generated randomly.

We ran experiments with the heuristic, starting with a 27 snake-shaped network that is fully optimized. The results are depicted in Figures 7 and 8. In each figure there are three lines. The gray line (connected with squares) represents the optimization results of Schema's simulator, the black line (connected with triangles) represents the original distributed algorithm described in this paper, and the blue line (connected with circles) represents the heuristic results.

Using the heuristic at a density of 10, the percentage of served mobiles increases from 78.36% (in the distributed algorithm, without the heuristic) to 78.56% on average (an improvement of 1% in the number of non-served users). In comparison to the regular mechanism used earlier ("Schema's results"), the percentages of served users increases from 72.57% to 78.56% on average (an improvement of 21.8% in the number of non-served users). At a density of 4 the percentage of served mobiles increases from 96.70% (in the distributed algorithm, without the heuristic) to 97.13% on average (an improvement of 13% in the number of non-served users). In comparison to the regular mechanism used earlier, the percentages of served users significantly improves and increases from 96.32% to 97.13% on average (an improvement of 22% in the number of non-served users).

It is interesting to note, that Schema's results are very stable, yet less efficient than our algorithm. The heuristic is more efficient than our original algorithm, however less stable. Thus, there is a tradeoff between the stability of the different algorithms in terms of the mobiles served, and the ability to improve the overall network performance in the short term in environments with changing traffic.

6. CONCLUSIONS

This paper demonstrates the promise embodied in local agent-based methods to improve overall network performance. The integration of a distributed local negotiation mechanism for radio network simulations, can affect the real-time adaptation of deployed networks. We have shown the applicability of our proposed method using two distinct parameters, scenarios and baselines, which bolsters our confidence with regard to its efficacy for general scenarios. The introduction of an efficient preliminary heuristic to produce the negotiated offers themselves improves the performance of the mechanism. Finding more heuristics and incorporating the changes applied to a set of parameters as part of each local negotiation, are expected to significantly increase the magnitude of improvement in future generation networks.

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

Future work in this field includes the introduction of clusters and intra-cluster negotiation, as well as the investigation of the nature of 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. Other heuristics can also be employed in order to gain an increase in overall network performance.

7. 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, Haim Pistner and Hezi Sayegh for their help in understanding the infrastructure of the cellular network and 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.

8. REFERENCES

[1] Actix. http://www.actix.com/products/acp.htm.

[2] Andrew. http://www.andrew.com/products/measurement_sys/odyssey/.

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

[4] H. Bölcskei and A. J. Paulraj. *The Communications Handbook*, chapter Multiple-input multiple-output (MIMO) wireless systems, pages 90.1 – 90.14. CRC Press, second edition, 2002.

[5] S. C. Borst, A. Buvaneswari, L. M. Drabeck, M. J. Flanagan, J. M. Graybeal, G. K. Hampel, M. Haner, W. M. MacDonald, P. A. Polakos, G. Rittenhouse, I. Saniee, A. Weiss, and P. A. Whiting. A dynamic optimization in future cellular networks. *Bell Labs Technical Journal*, 10(2):99–119, 2005.

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

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

[8] F.Ikegami, T.Takeuchi, and S.Yoshida. Theoretical prediction of mean field strength for urban mobile radio. *IEEE Transactions on Antennas and Propagation*, 39(3):299–302, March 1991.

[9] S. Frattasi, H. Fathi, F. Fitzek, M. Katz, and R. Prasad. Defining 4g technology from the user perspective. *IEEE Network Magazine*, 20(1):35–41, January 2006.

[10] J.Walfisch and H. Bertoni. A theoretical model of uhf propagation in urban environments. *IEEE Transactions on Antennas and Propagation*, 36:1788–1796, 1988.

[11] R. Lin, D. Dor-Shifer, S. Kraus, and D. Sarne. Local negotiation in cellular networks: From theory to practice. In *The Eighteenth Conference on Innovative Applications of Artificial Intelligence (IAAI)*, 2006.

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

[13] A. Raniwala and T. cker Chiueh. Architecture and algorithms for an ieee 802.11-based multi-channel wireless mesh network. In *IEEE Infocom*, pages 2223–2234, 2005.

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

[15] S. Saunders. *Antennas and Propagation for Wireless Communication Systems*. John Wiley, UK, 1999.

[16] Schema. http://www.schema.com/products_umts.html, 2005.

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

[18] I. Siomina. P-cpich power and antenna tilt optimization in umts networks. In *AICT/SAPIR/ELETE 2005*, pages 268–273, July 2005.

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

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