ARTICLE IN PRESS

Ad Hoc Networks xxx (2011) xxx–xxx

Contents lists available at ScienceDirect

Ad Hoc Networks

journal homepage: www.elsevier.com/locate/adhoc

Mitigating multi-path fading in a mobile mesh network

Marcos A.M. Vieira *, Matthew E. Taylor, Prateek Tandon, Manish Jain, Ramesh Govindan, Gaurav S. Sukhatme, Milind Tambe

Department of Computer Science, University of Southern California, Los Angeles, CA, United States

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Mobile and wireless ad-hoc networks
Multi-path fading
Complex environments
Distributed optimization
Experimental and prototype results

ABSTRACT

By using robots as routers, a team of networked robots can provide a communication substrate to establish a wireless mesh network. The mobile mesh network can autonomously optimize its configuration, increasing performance. One of the main sources of radio signal fading in such a network is multi-path propagation, which can be mitigated by moving the senders or the receivers on the distance of the order of a wavelength. In this paper, we measure the performance gain when robots are allowed to make such small movements and find that it may be as much as 270%. Our main contribution is the design of a system that allows robots to cooperate and improve the real-world network throughput via a practical solution. We model the problem of which robots to move as a distributed constraint optimization problem (DCOP). Our study includes four local metrics to estimate global throughput.

1. Introduction

With advances in processor, memory, sensing, actuation, and radio technology, it is possible to assemble novel systems using off-the-shelf components. A good example is a robot with navigation capabilities, and an on-board processor with wireless communication capabilities. Among its many uses, such a robot can be used as a *router* in a mobile wireless mesh network. In such a network, a team of robots can provide a communication substrate for a collection of clients.

Such a mobile wireless mesh network can have applications in various settings. In an infrastructure-less settings, it can be used to form a connection backbone, such as in the LANdroids project [1], where the goal is enable soldiers to communicate even in dense urban settings. A mobile wireless network network can also be quickly and autonomously deployed in urban search and rescue efforts, allowing searchers to communicate even when no other infrastructure exists: thus, small robots could venture where humans cannot, to search for survivors of earthquakes, collapsed mines and other disasters.

Unlike a static, manually deployed mesh network, the dynamism in the network allows nodes to move and reorganize the network, to achieve optimize or improve coverage, performance, or other such objectives. In this paper, we leverage this mobility to consider a specific kind of performance improvement. Our work is motivated the observation that one of the main sources of radio signal fading in urban settings is multi-path propagation. Multi-path occurs when a transmitted signal takes more than one path to a receiver, causing the signals to interfere. The central observation of our paper is that robots can *actively* reduce multi-path effects by making small movements (or *micro-motion*). By avoiding deep fades, robotics routers can increase network throughput, enabling applications with higher bandwidth requirements, or improving user satisfaction in general.

Thus, in this paper we explore two questions. First, is it possible to improve the overall network throughput of a mobile mesh network by using (possibly coordinated) robotic micro-motion? Second, how would one design an on-line system that performed this optimization autonomously?

* Corresponding author. Tel.: +55 31 3409 5860.
E-mail addresses: mvieira@usc.edu (M.A.M. Vieira), taylorm@usc.edu (M.E. Taylor), tandonp@usc.edu (P. Tandon), manishja@usc.edu (M. Jain), ramesh@usc.edu (R. Govindan), gaurav@usc.edu (G.S. Sukhatme), tambe@usc.edu (M. Tambe).

1570-8705/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.adhoc.2011.01.014

In our work, we make minimal assumptions about the capability of robots and the available information. We consider the scenario where the robots do not need any knowledge of the environment. Robots do not have a map of the physical layout or known obstacle points. They also have no *a priori* topological information about wireless connectivity or interference maps. It is often difficult to predict signal propagation characteristics within an urban environment, since many factors can impact signal strength such as the angle of incidence, emitter location, and even the building materials. We do not use models of radio propagation but instead, we rely on empirical on-line measurements to make decisions about when to move. Finally, in our setting, robots are not constantly in motion, but only execute small movements relative to their neighbors: this is appropriate, given our goal is to explore how effective such motions are in improving performance. As a result, our work does not require robot localization or sophisticated navigation capabilities.

1.1. Contributions

Our paper makes four main contributions. First, we show experimentally that we can obtain up to a factor of 2.7 improvement in TCP throughput on our testbed via robotic micro-motion. This is encouraging, since our experiments were fairly adversarial, suggesting that similar, or even higher, gains could be achieved in other environments. Second, we present the design of a practical system for coordinated robotic micro-motion. This system contains a novel use of a distributed constraint optimization framework: in this framework, robots make local measurements of a wireless performance metric, then decide, in coordinated fashion, which robot should move, and in what direction. This computation is executed iteratively, until the network converges to an improved throughput state. An important component of this framework is the choice of wireless performance metric: we empirically explore four different metrics, and show that a carefully chosen local metric can achieve a near-optimal performance. Finally, we evaluate our system with physical robots in an indoor environment and demonstrate that we are able to achieve an average global throughput improvement of 30% while maximizing only local metrics and with no *a priori* knowledge of the environment. Prior research [2,3] showed that, for a pair of nodes, micro-motion can increase receive signal power and improve packet reception more than any coding scheme could achieve. We, on the other, focus on an approach to improve global network throughput using explicit coordinated micro-motion. To the best of our knowledge, no prior research has designed a practical system to take advantage of antenna gain resulting from micro-motion, nor has any work explore what throughput improvements are achievable.

Our paper is structured as follows. In the next section, we start by providing the requisite background, and motivating the problem setting (Section 2). In Section 3, we firstly describe our robot platform. We then validate that network performance can improve from the robotic micro-motion by measuring TCP and UDP Throughput. Then, we describe the distributed constraint optimization framework (Section 4) and describe the design of our system. Finally, we describe our experimental methodology, and our main results in Section 5.

2. Background

Radio signal fading can be attributed to two mutually independent phenomena: multi-path propagation and path loss. We briefly discuss these phenomena and further details can be found elsewhere [4].

Multi-path propagation is a small-scale effect where the distance scales involved are on order of a wavelength. Multi-path occurs when a transmitted signal takes more than one path to a receiver, causing the signals to interfere. Interference has either a constructive or destructive effect on the main component depending on whether it arrives in or out of phase.

Path loss is a large-scale effect of propagation in any medium (e.g., air or water), defined by the way in which radio energy is transmitted in the medium of propagation and its resulting loss. This property is also called *slow fading*.

In this paper, we focus on mitigating the destructive interference arising from multi-path fading. Two models of multi-path fading have been described in the literature: as described in [5], if all signal components that reach the receiver are of equal strength, the multi-path fading is called Rayleigh fading, while if there is a line-of-sight (LoS) component that is significantly stronger, we have Ricean fading. In either case, small movements of the radio can help mitigate the effect of deep fades (strong destructive interference), and this is the observation we experimentally explore in this paper. Lindhé et al. [5] use Rayleigh fading model with data correlated until 0.38 wavelengths, while others [1] have suggested that moving $\frac{1}{4}$ to $\frac{1}{2}$ of a wavelength (λ) is sufficient to escape a deep fade. However, to our knowledge, we are the first to quantify TCP-level throughput improvements resulting from micro-motions, and also the first to design a practical decentralized coordination strategy to exploit micro-motions to obtain performance improvements.

In our work, we move our radios on the order of half wavelength so that signals in different locations are uncorrelated, helping our nodes escape deep fades. However, such movements cannot be performed independently. Fading for one radio is defined with respect to a single neighbor: a local movement may allow the radio to escape one deep fade, but at the same time introduce a new fade with respect to a different neighbor. Thus, it is critical to coordinate movements to improve the overall throughput. In this paper, we address a series of questions: (1) Is there a sequence of coordinated movements that improves the throughput? (2) How well will using only local information allow us to optimize the network (relative to the global optimum configuration)?

3. Can micro-motion improve throughput?

In this section, we show experimentally that TCP and UDP throughput can be improved by a much as a factor of 2.7 via robotic micro-motion in our testbed.

Please cite this article in press as: M.A.M. Vieira et al., Mitigating multi-path fading in a mobile mesh network, Ad Hoc Netw. (2011), doi:10.1016/j.adhoc.2011.01.014

To investigate the efficacy of small movements in improving mesh network throughput, we have used physical robots and conducted experiments in an office building. This section describes our platform and our experimental methodology, and then presents the results.

3.1. The robot platform

We use a commoditized robotics platform and made minimal modifications to it using commercial off-the-shelf products. Our platform consists of an iRobot Create and a small embedded computer mounted on top of it (Fig. 1).

The Create, a differential drive robot, has a round chassis with a diameter of 33 centimeters. The robot has two kinds of sensors.

First, it has a pair of tactile sensors that, together with a bumper, can help determine if the robot hits an obstacle and the angle at which it does so. Second, it has a suite of infrared (IR) sensors: the bumper contains an IR wall sensor on the right and an omnidirectional IR receiver in the top, and four additional IR sensors mounted underneath the bumper facing down. We do not add additional sensing hardware to the Create.

The embedded computer, the Ebox 3854, is an 800 MHz embedded PC with 256 MB shared DDR memory, and supports a 1280 × 1024 VGA interface, one 10/100 LAN, and USB, mini PCI and compact flash sockets. The embedded computer runs Ubuntu (Linux Kernel 2.6.22) as the operating system. For sensing and control, we developed a Create driver for Player [6]. Robot navigation incurs errors in odometry over larger distances. But, given that in our framework, the distances are very small (at most 6 cm), the navigation error is very small. For the speed of 0.20 cm/s, we measured less than 1cm translation error and less than 0.14 radian rotation error (as measured with a multifunction knob Griffin PowerMate attached to the robot).

3.2. Configuration

Our initial experiments use five robots distributed throughout an indoor office environment as shown in Fig. 2. Robots 1, 2 and 3 are within line-of-sight of each other, and the other robots are each not within line-of-sight of any robot. Fig. 1 depicts part of our experiment, showing robots 1–3.

Fig. 1. This picture shows part of the experiment set up, which has a team of iRobot Creates with an Ebox.

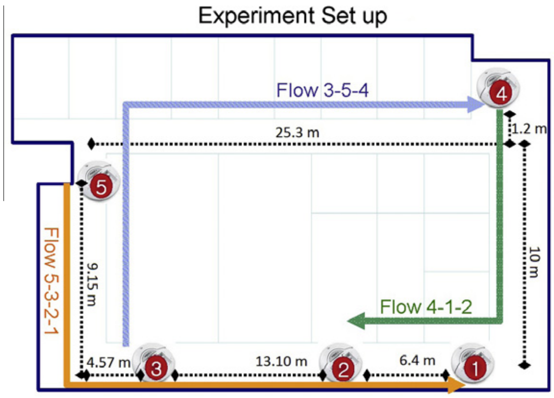

Fig. 2. Initial configuration of team of Creates.

The robots were configured to use 802.11b, with an 11 Mbits/s data rate (maximum data rate), in ad-hoc mode. The transmission power was set to the lowest possible value so we could experiment with as many robots as possible. We use channel 14 (which is unused by commercial cards in the US), ensuring that *we do not observe external interference*. The network was configured with *static routing* to avoid routing flapping (a router forwards packets via one route then changes to another router) interfering with the measurements.

The network has three multi-hop flows, represented by arrows in Fig. 2. The flows go through nodes 3-5-4, 4-1-2 and 5-3-2-1, which will be referred to as the 3–4 flow, 4–2 flow and 5–1 flow respectively. The flows use all the links in the network. Each flow takes 10 s for each sampling. *We avoid interference between flows* by starting and measuring the flows sequentially. This is deliberate: our objective is to determine if micro-motion can enable good path selection overall and running simultaneous flows would not have allowed us to observe the impact of improved throughput as a result of improving link quality because the simultaneous flows would interfere with each other. Flows were created with the *iperf* [7] tool and SNR values were measured (per link) using the *iwspy* Linux utility.

3.3. Throughput improvement

The first set of experiments quantifies the throughput improvement obtainable from small movements. The basic methodology is to exhaustively evaluate all possible configurations resulting from each robot executing a micro-motion, and then measuring the throughput achieved by all the flows in each configuration. Since the total number of possible configurations is exponential with the number of possible robot locations, we constrain robots to only two positions for tractability. In a later section, we will relax this assumption, allowing for more positions per robot.

We use five robots in our experiment, yielding a total of 32 possible configurations where each configuration is the average sampled of at least five times. During all the measurements, the noise was constant at −98 dBm, which implies there were no external time-varying radio sources contributing signal interference.

In this section, we present results for both TCP and UDP. We consider TCP flows because TCP is the most commonly used transport protocol. For completeness, we also present results for UDP flows. The throughput loss for TCP and UDP are within a constant factor of each other, suggesting that the performance loss comes from packet drops as a result of poor link quality, and not any other TCP artifact.

Figs. 3 and 4 show the multi-hop TCP and UDP flow, respectively, per configuration. Indeed, there is high variability between the flows per configuration in both TCP and UDP flows. Flow 3–4 is the flow with highest variance. Flow 5–1 is the flow with the longest hops and has lowest throughput on average. Fig. 3 shows that flow 3–4 has significant differences in TCP throughput. We can also conclude that some flows will improve while others will degrade such as flows 4–2 and 5–1 for configurations 16 and 20.

To quantify the variance, we sort the sum of the throughputs of all the TCP (respectively UDP) flows in each configuration. Fig. 5 (and Fig. 6) shows that there is a

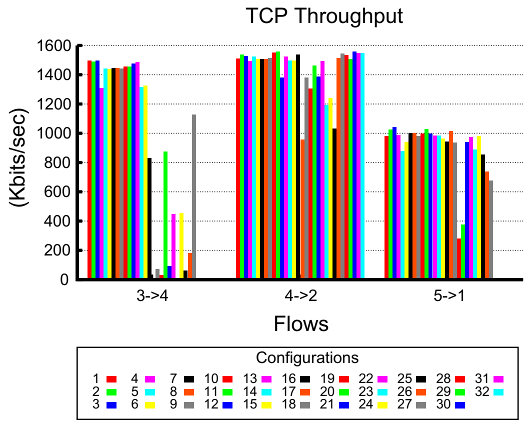

Fig. 3. TCP flows.

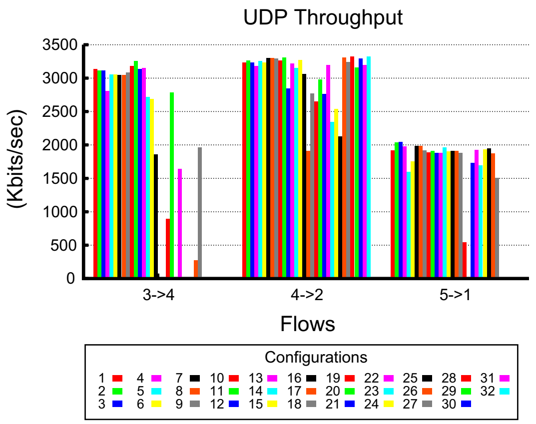

Fig. 4. UDP flows.

significant difference across configurations (recall that each configuration can be attained from a starting configuration by micro-motions of a subset of the five robots). There is almost a *2.5× difference* in total throughputs between the best configuration and the worst. Our topology has not been especially engineered to achieve this result, which leads us to believe that in other topologies we are likely to see similar performance improvements. This suggests that a mechanism for coordinated small movements can improve performance significantly.

Now, we illustrate the variations in each link per position. Fig. 7a–e show the SNR per link for each node for all possible configurations. We can also visualize the variance of SNR for each configuration. As expected, nodes with non line-of-sight connectivity have lower SNR. We can conclude that one-hop metric SNR varies from micro-motions.

From this section, we conclude that there is a significant difference in UDP and TCP throughput obtainable from micro-motions. In the next section, we discuss how to design a system to take advantage of micro-motions to improve throughput.

4. Using distributed reasoning for micro-motion based throughput improvement

In this section, we describe the distributed constraint optimization framework and how we use it to design a decentralized method for throughput improvement in mobile mesh networks.

A distributed constraint optimization problem (DCOP) consists of a set V of n variables, $\{x_1,x_2,\ldots,x_n\}$, assigned to a set of agents (e.g., independent reasoning entities), where each agent controls one variable's assignment. Variable x_i can take on any value from the discrete finite domain D_i. The goal is to choose values for the variables such that the sum over a set of binary constraints and associated payoff or reward functions, f_{ij}: $D_i \times D_j \rightarrow N$, is maximized. More specifically, to find an assignment, A, s.t. $F(A)$ is maximized: $F(A) = \sum_{x_i,x_j \in V} f_{ij}(d_i,d_j)$, where $d_i \in D_i$, $d_j \in D_j$ and $x_i \leftarrow d_i$, $x_j \leftarrow d_j \in A$. For example, in Fig. 8, x_1, x_2, and x_3 are variables, each with a domain of $\{0,1\}$ and the reward function as shown. If agents 2 and three choose the value 1, the agent pair gets a reward of 9. If agent 1 now chooses value 1 as well, the total solution quality of this complete assignment is 12, which is locally-optimal as no single agent can change its value to improve its own reward (and that of the entire DCOP). $F((x_1 \leftarrow 0),(x_2 \leftarrow 0),(x_3 \leftarrow 0)) = 22$ and is globally optimal.

In this problem, we model each mobile radio as an agent. Every value an agent can take is one possible physical position for the mobile radio. Constraint exist between neighbors in the wireless network. Rewards on the constraints are defined by a local metric, such as the packet reception rate on the wireless link between two neighbors.

While there are many approaches to solving DCOPs, we implemented the Maximum Gain Message (MGM [8]) DCOP method. The MGM algorithm will find a locally-optimal assignment of values for all agents. MGM defines a round as a period in which every agent: (1) communicates its current value to all its neighbors, (2) calculates and

Please cite this article in press as: M.A.M. Vieira et al., Mitigating multi-path fading in a mobile mesh network, Ad Hoc Netw. (2011), doi:10.1016/j.adhoc.2011.01.014

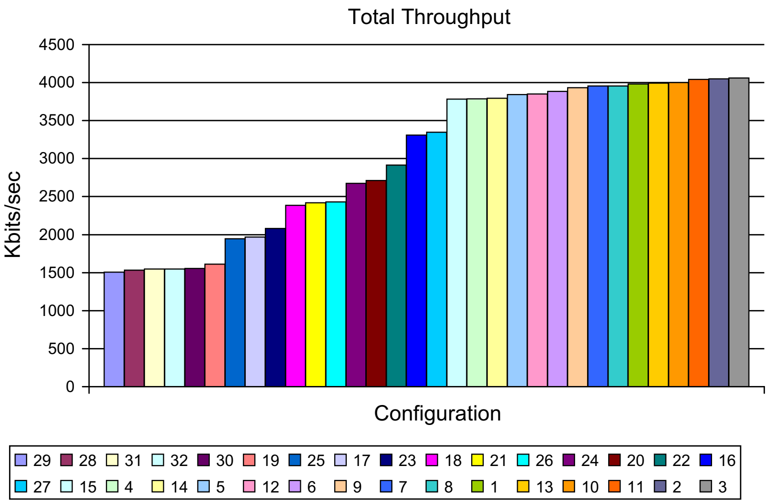

Fig. 5. Sorted total TCP flows.

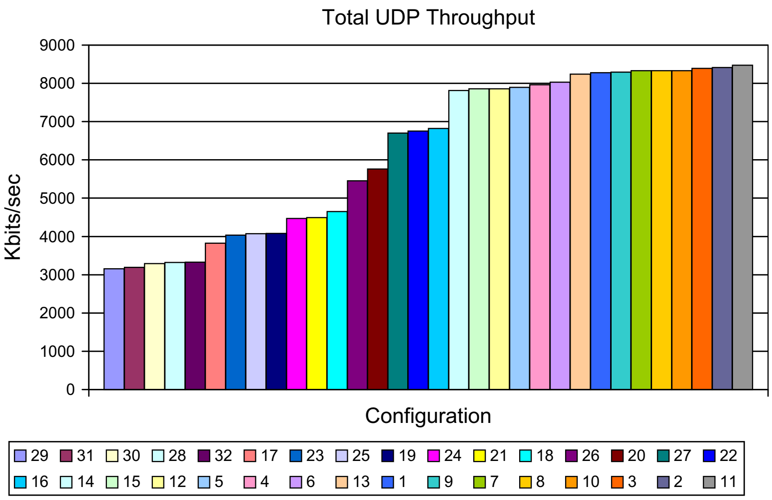

Fig. 6. Sorted total UDP flows.

communicates its *bid* (the maximum gain in its local reward if it is allowed to change values) to all its neighbors, and (3) changes its value (if allowed). An agent is allowed to change its value if its *bid* is larger than all the bids it receives from its neighbors. At quiescence, no one agent can deviate from the proposed assignment and increase the net reward. We denominate MGM-Omniscient when the agents have the reward for each possible value. MGM-Omniscient gives an upper bound.

The agents in a DCOP are traditionally assumed to have *a priori* knowledge of the corresponding reward functions. In order to more flexibly model a class of real world domains, we previously introduced *Distributed Cooperative Exploration and Exploitation* (D-CEE) [9] problems, which do not make this assumption. Thus, D-CEE problems appear similar to DCOPs, but with the following features absent from DCOPs: (1) agents initially know the constraint graph but only discover rewards when a pair of agents set their values to explicitly discover a reward value, (2) problems last a set amount of time, and (3) the agents' seek to maximize the on-line global reward over this time horizon *T*.

The mapping from our network optimization problem onto a D-CEE is similar to that of a DCOP, with one important difference. Agents (robots in our case) must explore different locations to determine the value of local (point-to-point) metrics, and we provide a time horizon after which the agents must stop optimizing (to ensure that the network converges quickly) and the on-line reward is maximized (ensuring that the network will quickly improve, and that it will be performing as well as possible during the optimization).

SE-Mean [9], a D-CEE algorithm used in this paper, assumes the average reward (denoted μ) on each constraint for all unexplored values for agents. On every round, each agent bids its expected gain: $NumberLinks \times \mu - R_c$ where

Please cite this article in press as: M.A.M. Vieira et al., Mitigating multi-path fading in a mobile mesh network, Ad Hoc Netw. (2011), doi:10.1016/j.adhoc.2011.01.014

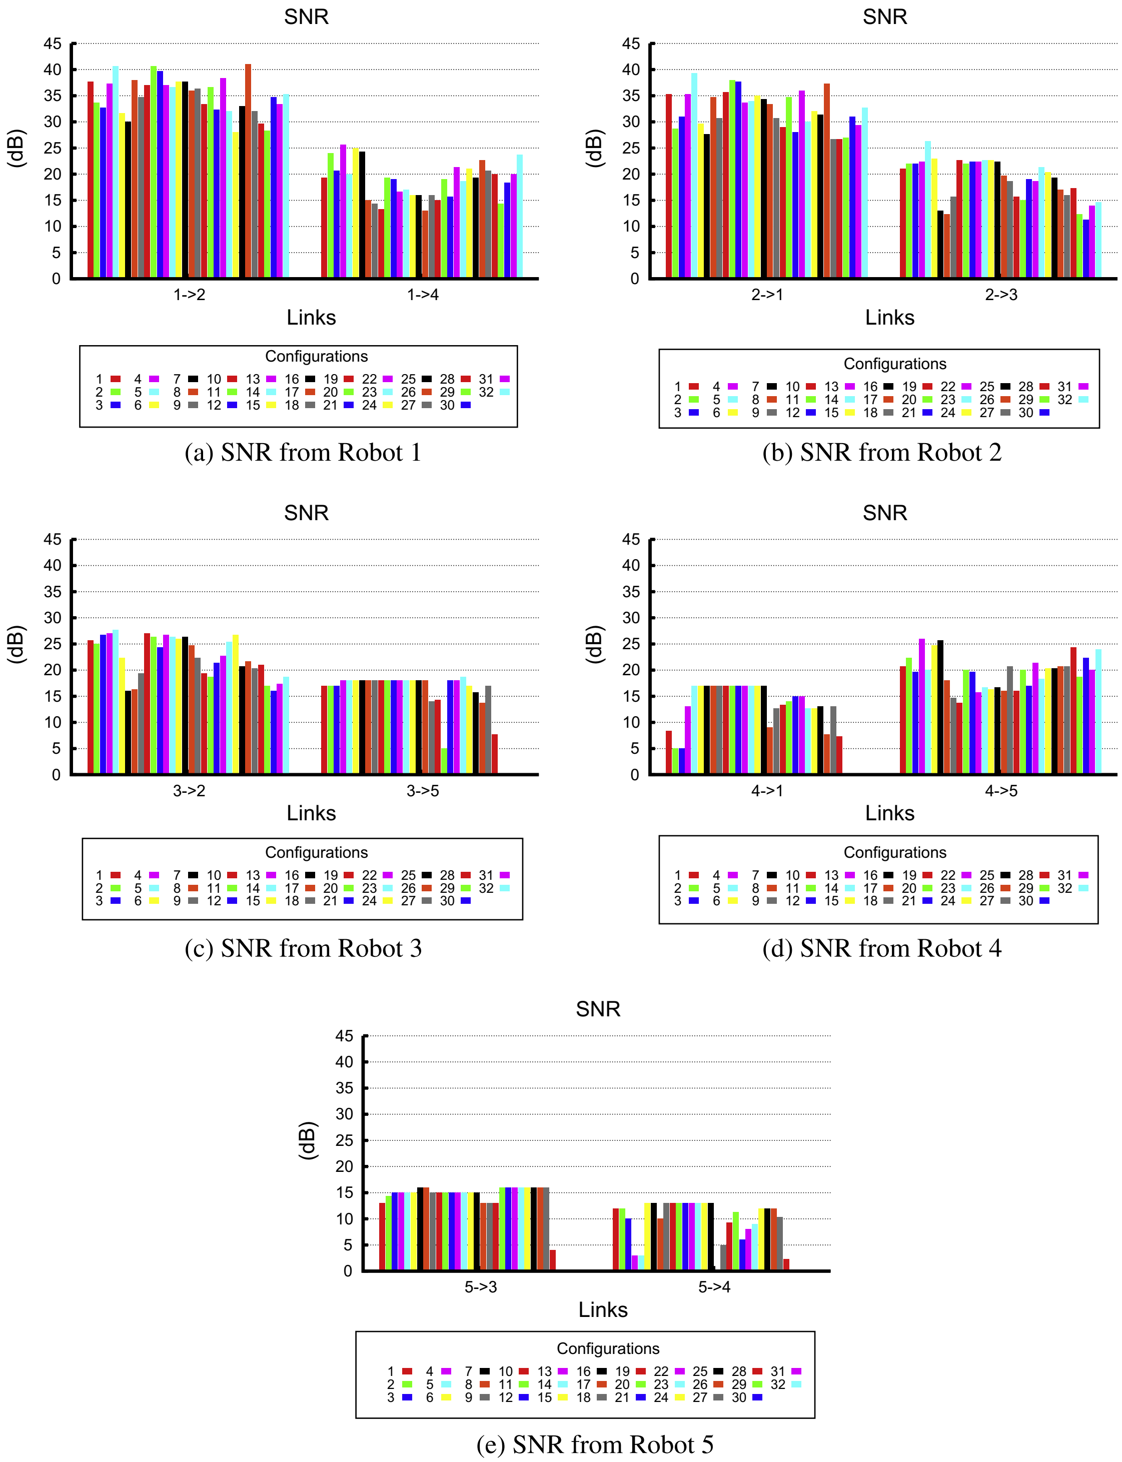

Fig. 7. SNR for all configurations.

R_c is the current reward. The algorithm then proceeds as in MGM-Omniscient. This algorithm causes the agents to greedily explore until they achieve the average reward, allowing them to converge on an assignment.

The overall algorithm thus consists of two phases. In the first phase, each robot independently, and without coordinating with other robots, samples the local metric. After computing the local metric with respect each neighbor,

Please cite this article in press as: M.A.M. Vieira et al., Mitigating multi-path fading in a mobile mesh network, Ad Hoc Netw. (2011), doi:10.1016/j.adhoc.2011.01.014

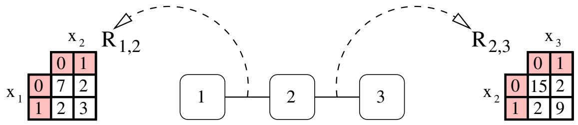

Fig. 8. This figure depicts a three agent DCOP.

each agent calculates the mean (μ) and standard deviation to be used as input to the SE-Mean. The second phase consists of running the DCOP as explained above. The pseudocode of the algorithm is presented in Algorithm 1.

Algorithm 1. D-CEE Algorithm

```
{Explore Phase}
for subset of positions do
    Sample
    Move
end for
Calculate average reward μ
{Second Phase – Extended MGM Algorithm}
repeat
Sample
Communicate current value
Calculate and Communicate bid
Winner of bid Moves
until achieve average reward over all neighbors
```

In summary, we model the problem of maximizing throughput as a DCOP/D-CEE problem, where robots must coordinate their movements in a decentralized fashion. The overall goal is to maximize the throughput of a set W of (possibly multi-hop) flows w_{ij} between nodes, but by using purely local metrics. These local metrics enable nodes to effectively use micro-motion to escape from deep fades, improving link quality as well as w_{ij}. An important component of our design is the choice of the appropriate local metric. As it turns out, this choice makes a significant difference to the performance of our algorithm, which we evaluate in the next section.

5. Results

In this section, we experimentally evaluate our system using the distributed constraint optimization framework. Initially, we explore four different local metrics, and show that the choice of local metric is important for the performance of our algorithm. Thereafter, we quantify the performance improvement attainable in practical settings by demonstrating end-to-end evaluations of our algorithm.

5.1. Local metrics

This section evaluates the possible local metrics to be used as a local reward to the coordination algorithms so that the system can achieve a global reward improvement. These results will prove essential to understanding the system performance, described in the following subsection.

We evaluate four local metrics, each of which defines the reward of a particular agent (and thus how likely it will attempt to change its position):

- *minimum SNR* (Signal-to-Noise Ratio) is the minimum SNR on an agent's links
- the *summation of SNR* is the sum of all SNRs on an agent's links
- the *minimum PRR* (Packet Reception Rate) is the minimum PRR on an agent's links
- and the *summation of PRR* is the sum of all PRRs on an agent's links

Our specific choice of these four metrics is driven by their simplicity: these metrics can be estimated cheaply and quickly so that network reconfiguration can be done faster than if other, more heavyweight methods (such as direct throughput measurement) were used. Our work borrows heavily from the wireless literature, which has long used SNR [10–12] and PRR [13–15] as predictors for link quality and throughput in 802.11 radios.

However, a link quality metric alone does not define a reward function. In general, each node in a network may have many neighbors, and the reward function is defined per node. There are two natural choices for the reward function for a node: the *min* of the link quality metric (either SNR or PRR) over all neighbors, or their *sum*. This results in four choices for the local metric, which we evaluate below.

To evaluate the local metrics, we conducted experiments with the same configuration as in Section 3.2. There are five robots, each with two possible positions, yielding a total of 32 possible configurations where each configuration was sampled at least five times. We present results for TCP flows because TCP is the most commonly used transport protocol. As shown in Section 3 the results for TCP and UDP were very similar, modulo a scale factor.

We have the ground truth for the experiments as we measured the actual throughput for each flow in every configuration. In this way, we know which configuration was optimal. It is also important to notice that we do not need to estimate (SE-Mean) the reward matrices since we collected data for all possible configurations.

We evaluate the overall system improvement obtained by using MGM-Omniscient. We focus on the four local metrics. We evaluated each local metric using the data obtained from exhaustive search experiments. Recall that each agent will work to maximize its local reward (in this case, the sum or max of the SNR or PRR on its links), which will ideally maximize the global metric. Although the agents work to maximize SNR and PRR, this section shows that the corresponding network flows are also maximized, *even though they are not directly measured* by the agents for optimization purposes.

We compare the best local metric to predict global gain. Fig. 9 shows how close to the optimal the configurations are when using the local metric. First, we can conclude that the local metric matters when designing the system. For instance, the min SNR metric improves the total throughput on average by almost 45%. The improvement is not higher than this because our approach uses a local metric to maximize a global metric, and because the local metric

Please cite this article in press as: M.A.M. Vieira et al., Mitigating multi-path fading in a mobile mesh network, Ad Hoc Netw. (2011), doi:10.1016/j.adhoc.2011.01.014

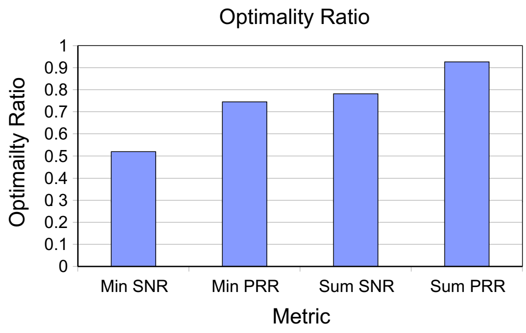

Fig. 9. Improvement per metric.

might incur prediction error. Secondly, we can also conclude that it is possible to achieve a near-optimal performance with a carefully chosen metric: as the figure shows, Sum PRR metric shows that it is possible to achieve the near-optimal performance.

In summary, we can improve throughput by carefully choosing a local metric (Sum PRR) within the DCOP Framework.

5.2. Optimizing with more positions

We have also conducted experiments in which robots are allowed to sample more positions. Our approach is generic and does not need to be modified to support additional positions. For our experiments, we allowed each robot to have five possible positions as illustrated in Fig. 10a. The work by Lindhé et al. [5] solves the problem of how many samples are needed for given communication performance.

It is not practical for agents to visit all positions since the number of configurations is exponential. Thus, in this case (unlike our experiments above where there were fewer positions) agents need to estimate the local metrics for unexplored positions, which is calculated by the Static Estimation (SE) Mean Algorithm.

Thus, the overall algorithm consists of two phases. In the first phase, each robot independently samples the local metric. There is no coordination and the robots sample simultaneously. After collecting the local metric to all its neighbors, agents calculate the mean (μ) and standard deviation to be used as input to the SE-Mean. The second phase consists of running DCOP as explained previously.

Fig. 10b shows the experiment's physical configuration. We run the experiments with five robots and three flows. Each robot can move to five positions. The SE-Mean algorithm is used to estimate the local metric values for the unexplored positions. Robot 2 has no free line-of-sight with respect to the others.

Fig. 10c illustrates the overall percentage gain over each round (line 8 of Algorithm 1). The base is the initial total throughput. The system gain is about 32%. Table 1 shows the gain and the number of rounds for three instances of

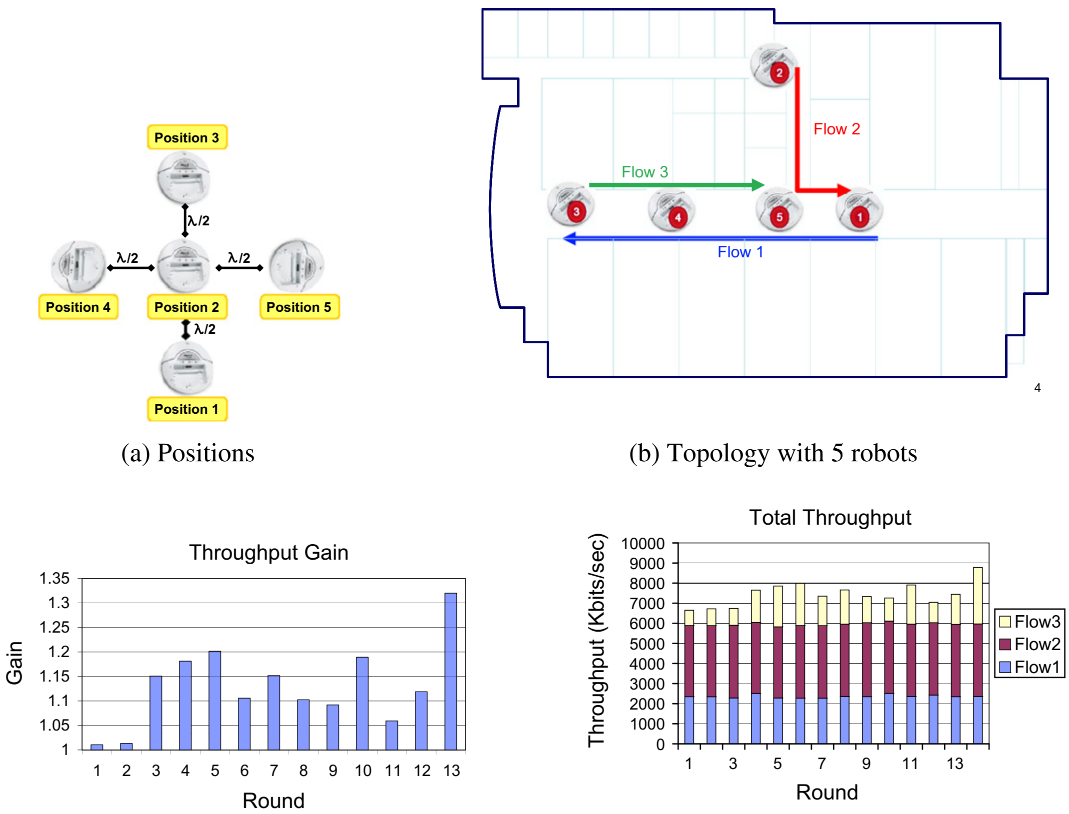

(a) Positions (b) Topology with 5 robots
(c) Throughput Improvement per Round (d) Throughput per round per flow

Fig. 10. Experiment results.

Table 1
Experiments with five robots.

Experiment	Gain	Rounds
#1	1.32	13
#2	1.27	9
#3	1.30	10

the experiments. In Section 3, we showed that up to a 2.5× performance improvement is possible with micro-motion. Our results from this experiment do not contradict our earlier finding: that finding discusses the performance difference between the worst and the best configuration. In this experiment, we started with an arbitrary configuration that was not guaranteed to be the worst and therefore the actual performance improvement we observe is not as high.

In some rounds, the overall gain decreased when compared to previous round. This is because the robots do not know the direction of its neighbors. When a robot moves, the robot might go to the opposite direction of the neighbor with the weakest local metric. The system can overcome this in the next round since robots will move again. The iterations stop when the local metric has improved more than a threshold (line 13 of Algorithm 1). These results are encouraging because, even with relatively simple local algorithms, and using small movements, we could significantly improvement system throughput.

Fig. 10d shows each flow's throughput per round. Two flows have little variance over time, while one flow's throughput increases significantly as a result of micro-motion. Therefore, we can conclude that by using SE-Mean to estimate unexplored positions, the system can handle many positions and still improve throughput.

5.3. Temporal variation

Could our performance improvements have been explained by temporal variation of wireless signals? To test this, we disabled robot motion so we could measure how the wireless signal varies over time. Fig. 11 depicts the static total throughput variation over time. The maximum variation is about 5%, which fails to explain the 35% improvements seen in our experiments. Thus, system improvements do not arise simply from changes in the environment.

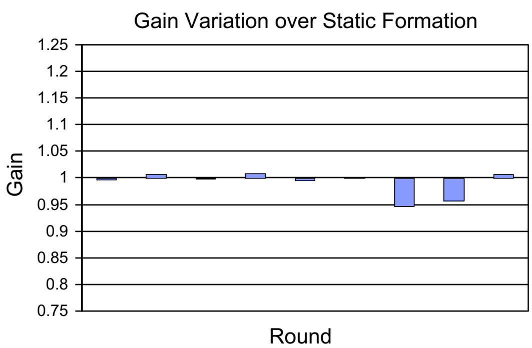

Fig. 11. Static variation.

5.4. Experiments with more robots

We also conducted experiments with more than five robots to validate our distributed system. Fig. 12a shows an experiment configuration which has seven robots and three multi-hop flows. Only robots 5, 6, and 7 have free line-of-sight with respect to each other. Each robot can move five positions. Table 2 shows the gain and the number of rounds for three instances of experiments. Fig. 12b illustrates the overall percentage gain over each round for a given experiment. In this configuration, we achieve an overall system-wide performance gain of about 30%, further validating our approach to distributed optimization via robot micro-motion.

6. Related work

Using small movements to combat the multi-path fading effects in complex environments has promise and this paper is not the first to examine such effects. In [2,3], the authors showed that, for a pair of nodes, micro-motion can increase receive signal power and improve packet reception more than any coding scheme could achieve. We, on the other, focus on an approach to improve global network throughput using explicit coordinated micro-motion. Other work includes [5], where the authors propose a methodology for exploiting multi-path fading by controlling the robot according to radio signal strength. They solve the problem of how many samples are needed for given communications performance and how they should be spaced and provide lower bounds on the number of samples for a single robot. Using 802.15.4 radio, they also show there is room for improvement (as much as 20 dB in RSSI).

Other approaches have leveraged more general forms of mobility (beyond micro-motion) for network throughput improvement or to build and configure mobile mesh networks. Early theoretical work [16] shows that mobility increases capacity with random source-destination pairs with loose delay constraints. Other work [17] considers the problem of controlling a team of robots to ensure end-to-end communication. To mitigate environmental interference, they propose two different metrics, point-to-point signal strength and data throughput, to monitor the network connectivity of the system. Even ad-hoc communication protocols pose difficult challenges during multi-robot experimentation, as shown by Zeiger et al. [18]. However, their focus is not on micro-motions, they need a map of the environment and optimizing network throughput is not one of their goals.

Complementary to our work [19], discusses a game-theoretic dynamic programming algorithm to guarantee that a single mobile user is connected to a base station by moving a chain of robotic routers.

Multiple-input multiple-output (MIMO) [20] techniques with multiple antennas [21] take advantage of spatial diversity and spatial multiplexing and can improve performance by avoiding deep fades through diversity. For example, consider the scenario of two transmit antennas at a node sending to one receiving antenna. This adds spatial diversity because of the independently faded paths.

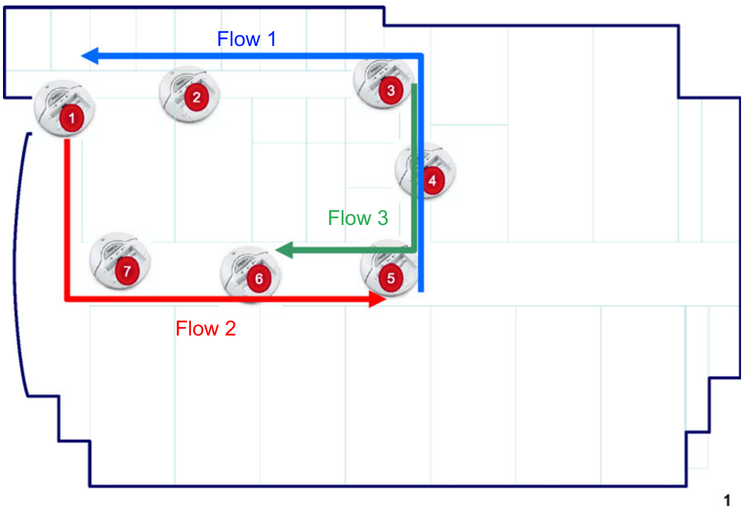

(a) Topology with 7 robots

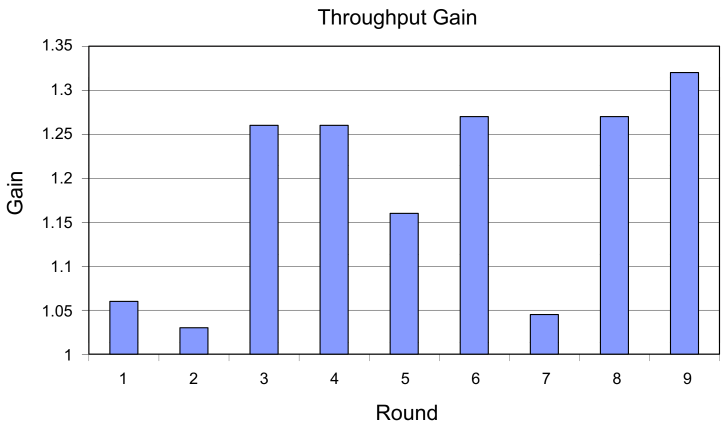

(b) Throughput Improvement per Round

Fig. 12. Experiment results.

Table 2
Experiments with seven robots.

Experiment	Gain	Rounds
#1	1.32	9
#2	1.28	9
#3	1.25	10

However, when node positions are fixed, there are limits to diversity gains. For example, in certain scenarios such as at low SNR, the extra transmit antennas make little difference in performance [21]. Our approach is complementary, since it uses explicit micro-motions to improve performance, and can improve performance in scenarios where MIMO gains are limited. We intend to investigate the performance gain of micro-motions with MIMO configuration in future work. For more information, we refer the interested reader to check textbooks on wireless communications and MIMO [22–24].

Delay-tolerant networking (DTN) [25,26] is a computer network that may lack continuous network connectivity but is still operable. DTNs can take advantage of mobility to deliver messages. Unlike DTNs, where nodes may only have intermittent connectivity, our work applies to mesh networks where a communication backbone exists in the network.

The distributed constraint optimization framework has been studied extensively in the multi-agent literature. In [9], the D-CEE framework is presented to study the problem of how to coordinate mobile nodes to maximize the cumulative RSSI. The paper's focus is on algorithms to study the trade-off between exploration and exploitation. We, on the other hand, focus on different local metrics (SNR, PRR) and how it affects the overall network. We quantify how much gain the network can benefit from small movements and how we can design a system to improve the real-world network throughput.

In addition to the DCOP work discussed in earlier sections, previous work in distributed constraint reasoning in sensor networks [27,28] uses a precursor method to the DCOP formulation which does not handle unknown reward matrices. Marder et al. [29] formulate dynamic sensor coverage as a "potential game," which is similar to a DCOP. However, like other DCOP work, the reward matrix is known, there is no time limit, and only final reward is considered. Cheng et al. [30] suggest an approach for coordinating a set of robots based on swarm intelligence, however the objective of the work is to disperse the robots evenly within a specified shape, and not to optimize the signal strengths across the network.

Correll et al. [31] look at optimizing a wireless network of mobile robots using a distributed swarm optimization, but are concerned with changing the topology (i.e., neighbors) of the network rather than optimizing signal strength. Gerkey et al. [32] address a similar problem, but use auction mechanism and the goals of agents are significantly different (agents modify the topology of the network and on-line reward is not emphasized). Farinelli et al. [33] perform decentralized coordination on physical hardware using factor graphs, however, rewards are known and cumulative reward is not considered.

7. Conclusion

In this paper, we demonstrate that mobile robots can be used successfully in a mesh network. With robotic routers forming a network, nodes can avoid deep face caused by multi-path fading. Our study shows that small movements can improve network performance and that the total network throughput could vary as much as 170% when the robots moved on the order of half a wavelength. Avoiding deep fades is a pairwise problem between sender and receiver, which requires coordination. Thus, we designed a practical system which uses the distributed constraint optimization framework to improve communication. We studied four local metric (min SNR, min PRR, sum SNR, sum PRR) to estimate the global throughput. Our results are encouraging because we can achieve an average global performance improvement of 30% while maximizing only local metrics.

References

[1] Broad Agency Announcement 07-46 LANdroids, <http://www.darpa.mil/ipto/programs/ld/ld.asp>.
[2] M. Lindhé, K.H. Johansson, Using robot mobility to exploit multipath fading – wireless communications in networked robotics, IEEE Wireless Communications 16 (1) (2009) 30–37.
[3] J. Smith, M. Olivieri, A. Lackpour, N. Hinnerschitz, RF-mobility gain: concept, measurement campaign, and exploitation, IEEE Wireless Communications 16 (1) (2009) 38–44.
[4] J.S. Seybold, Introduction to RF Propagation, Wiley-Interscience, 2005.
[5] M. Lindhé, K.H. Johansson, A. Bicchi, An experimental study of exploiting multipath fading for robot communications, in: Robotics: Science and Systems, Atlanta, GA, USA, 2007.
[6] B.P. Gerkey, R.T. Vaughan, K. Støy, A. Howard, G.S. Sukhatme, M.J. Mataric, Most valuable player: a robot device server for distributed control, in: International Conference on Intelligent Robots and Systems, vol. 3, Maui, HI, USA, 2001, pp. 1226–1231.
[7] A. Tirumala, F. Qin, J. Dugan, J. Ferguson, K. Gibbs, Iperf, <http://dast.nlanr.net/Projects/Iperf/>.
[8] J. Pearce, M. Tambe, Quality guarantees on k-optimal solutions for distributed constraint optimization, in: IJCAI '07: International Joint Conference on Artificial Intelligence, 2007.
[9] M. Jain, M.E. Taylor, M. Yokoo, M. Tambe, DCOPs meet the real world: exploring unknown reward matrices with applications to mobile sensor networks, in: IJCAI '09: Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence, 2009.
[10] K. Chebrolu, B. Raman, S. Sen, Long-distance 802.11b links: performance measurements and experience, in: MobiCom '06: Proceedings of the 12th Annual International Conference on Mobile Computing and Networking, ACM, New York, NY, USA, 2006, pp. 74–85.
[11] G. Judd,, X. Wang, P. Steenkiste, Efficient channel-aware rate adaptation in dynamic environments, in: MobiSys '08: Proceeding of the 6th International Conference on Mobile Systems, Applications, and Services, ACM, New York, NY, USA, 2008, pp. 118–131.
[12] D. Aguayo, J. Bicket, S. Biswas, G. Judd, R. Morris, Link-level measurements from an 802.11b mesh network, SIGCOMM Computer Communications Review 34 (4) (2004) 121–132.
[13] A. Jindal, K. Psounis, The achievable rate region of 802.11-scheduled multihop networks, IEEE/ACM Transactions on Networking 17 (4) (2009) 1118–1131.
[14] Y. Li, L. Qiu, Y. Zhang, R. Mahajan, E. Rozner, Predictable performance optimization for wireless networks, in: SIGCOMM '08: Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication, ACM, 2008, pp. 413–426.
[15] K.-H. Kim, K.G. Shin, On accurate measurement of link quality in multi-hop wireless mesh networks, in: MobiCom '06: Proceedings of the 12th Annual International Conference on Mobile Computing and Networking, ACM, 2006, pp. 38–49.
[16] M. Grossglauser, D.N.C. Tse, Mobility increases the capacity of ad hoc wireless networks, IEEE/ACM Transactions on Networking 10 (4) (2002) 477–486.
[17] M.A. Hsieh, A. Cowley, V. Kumar, C.J. Taylor, Towards the deployment of a mobile robot network with end-to-end performance guarantees, in: Proceedings of the 2006 IEEE International Conference on Robotics and Automation, 2006, pp. 2085–2091.
[18] F. Zeiger, N. Kraemer, K. Schilling, Commanding mobile robots via wireless ad-hoc networks – a comparison of four ad-hoc routing protocol implementations, in: ICRA '08: International Conference on Robotics and Automation., 2008, pp. 590–595.
[19] O. Tekdas, V. Isler, Robotic routers, in: ICRA '08: International Conference on Robotics and Automation, 2008.
[20] A.J. Paulraj, D. Gore, R.U. Nabar, H. Blcskei, An overview of MIMO communications: a key to gigabit wireless, in: Proceedings of IEEE, 2004, pp. 198–218.
[21] D. Halperin, W. Hu, A. Sheth, D. Wetherall, 802.11 with multiple antennas for dummies, SIGCOMM Computer Communications Review 40 (1) (2010) 19–25.
[22] A. Goldsmith, Wireless Communications, Cambridge University Press, New York, NY, USA, 2005.
[23] C. Oestges, B. Clerckx, MIMO Wireless Communications: From Real-World Propagation to Space-Time Code Design, Academic Press Inc., Orlando, FL, USA, 2007.
[24] D. Tse, P. Viswanath, Fundamentals of Wireless Communication, Cambridge University Press, New York, NY, USA, 2005.
[25] K. Fall, A delay-tolerant network architecture for challenged internets, in: SIGCOMM '03: Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, ACM, New York, NY, USA, 2003, pp. 27–34.
[26] S. Jain, K. Fall, R. Patra, Routing in a delay tolerant network, in: SIGCOMM '04: Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, ACM, New York, NY, USA, 2004, pp. 145–158.
[27] V. Lesser, C. Ortiz, M. Tambe, Distributed Sensor Nets: A Multiagent Perspective, Kluwer Academic Publishers, 2003.
[28] W. Zhang, Z. Xing, G. Wang, L. Wittenburg, An analysis and application of distributed constraint satisfaction and optimization algorithms in sensor networks, in: AAMAS '03: International Conference on Autonomous Agents and Multiagent Systems, 2003.
[29] J. Marden, G. Arslan, J. Shamma, Connections between cooperative control and potential games illustrated on the consensus problem, in: European Control Conference, 2007.
[30] J. Cheng, W. Cheng, R. Nagpal, Robust and self-repairing formation control for swarms of mobile agents, AAAI, 2005.
[31] N. Correll, J. Bachrach, D. Vickery, D. Rus, Ad-hoc wireless network coverage with networked robots that cannot localize, in: ICRA '09: International Conference on Robotics and Automation, 2009.
[32] B.P. Gerkey, R. Mailler, B. Morisset, Commbots: distributed control of mobile communication relays, in: AAAI Workshop on Auction Mechanisms for Robot Coordination, 2006.
[33] A. Farinelli, A. Rogers, A. Petcu, N. Jennings, Decentralized coordination of low-power embedded devices using the max-sum algorithm, in: AAMAS '08: International Conference on Autonomous Agents and Multiagent Systems, 2008.

Please cite this article in press as: M.A.M. Vieira et al., Mitigating multi-path fading in a mobile mesh network, Ad Hoc Netw. (2011), doi:10.1016/j.adhoc.2011.01.014

Marcos Augusto Menezes Vieira is an Assistant Professor of Computer Science at the Federal University of Minas Gerais (UFMG). He received his undergraduate and M.S. at the Federal University of Minas Gerais in Belo Horizonte, and M.S. and Ph.D. degrees in Computer Science from the University of Southern California (USC). He was the recipent of the CAPES fellowship. His research interests include Wireless Sensor Networks, Multi-Robot Systems and Distributed Embedded Systems.

Matthew E. Taylor is an assistant professor in the Computer Science Department at Lafayette College. He received his doctorate from the University of Texas at Austin in 2008, supervised by Peter Stone. After graduation, Taylor worked for two years as a postdoctoral research assistant at the University of Southern California with Milind Tambe. Current research interests include intelligent agents, multi-agent systems, reinforcement learning, and transfer learning.

Prateek Tandon received his undergraduate education at the University of Southern California (USC). His research interests include artificial intelligence, machine learning, robotics and their increasing synergy with computational neuroscience.

Manish Jain is currently a Ph.D. candidate at the University of Southern California. He is a part of the Teamcore Research group, led by Prof. Milind Tambe. His work is on modeling strategic interactions in multi-agent settings. He has received a Letter of Commendation from the city of Los Angeles for his contributions to the development of the security assistant deployed at LAX.

Ramesh Govindan received his B.Tech. degree from the Indian Institute of Technology at Madras, and his M.S. and Ph.D. degrees from the University of California at Berkeley. He is a Professor in the Computer Science Department at the University of Southern California. His research interests include scalable routing in internetworks, and wireless sensor networks.

Gaurav S. Sukhatme is a Professor of Computer Science (joint appointment in Electrical Engineering) at the University of Southern California (USC). He received his undergraduate education at IIT Bombay in Computer Science and Engineering, and M.S. and Ph.D. degrees in Computer Science from USC. He is the co-director of the USC Robotics Research Laboratory and the director of the USC Robotic Embedded Systems Laboratory which he founded in 2000. His research interests are in multi-robot systems and sensor/actuator networks. He has published extensively in these and related areas. Sukhatme has served as PI on numerous NSF, DARPA and NASA grants. He is a Co-PI on the Center for Embedded Networked Sensing (CENS), an NSF Science and Technology Center. He is a fellow of the IEEE and a recipient of the NSF CAREER award and the Okawa foundation research award. He is one of the founders of the Robotics: Science and Systems conference. He was program chair of the 2008 IEEE International Conference on Robotics and Automation and is program chair of the 2011 IEEE/RSJ International Conference on Robots and Systems. He is the Editor-in-Chief of Autonomous Robots and has served as Associate Editor of the IEEE Transactions on Robotics and Automation, the IEEE Transactions on Mobile Computing, and on the editorial board of IEEE Pervasive Computing.

Milind Tambe is a Professor of Computer Science and Industrial and Systems Engineering at the University of Southern California (USC). He leads the TEAMCORE Research Group at USC, with research is focused on agent-based and multi-agent systems. He is a fellow of AAAI (Association for Advancement of Artificial Intelligence) and recipient of the ACM (Association for Computing Machinery) "Autonomous Agents Research Award". He is also the recipient of the Christopher Columbus Fellowship Foundation Homeland security award, a special commendation given by the Los Angeles World Airports police from the city of Los Angeles, USC Viterbi School of Engineering use-inspired research award, Okawa foundation faculty research award, the RoboCup scientific challenge award, USC Steven B. Sample Teaching and Mentoring award and the ACM recognition of service award. He and his research group's papers have been selected as best papers or finalists for best papers at a dozen premier Artificial Intelligence and Operations Research Conferences and workshops, and their algorithms have been deployed for real-world use by several agencies including the LAX police, the Federal Air Marshals service and the Transportation security administration. He received his Ph.D. from the School of Computer Science at Carnegie Mellon University.

Please cite this article in press as: M.A.M. Vieira et al., Mitigating multi-path fading in a mobile mesh network, Ad Hoc Netw. (2011), doi:10.1016/j.adhoc.2011.01.014