

**Mobilized ad-hoc networks:
A reinforcement learning approach**

Yu-Han Chang
A.I. Lab, MIT
Cambridge, MA 02139
ychang@ai.mit.edu

Tracey Ho
LIDS, MIT
Cambridge, MA 02139
trace@mit.edu

Leslie Pack Kaelbling
A.I. Lab, MIT
Cambridge, MA 02139
lpk@ai.mit.edu

Abstract

Research in mobile ad-hoc networks has focused on situations in which nodes have no control over their movements. We investigate an important but overlooked domain in which nodes *do* have control over their movements. Reinforcement learning methods can be used to control both packet routing decisions and node mobility, dramatically improving the connectivity of the network. We first motivate the problem by presenting theoretical bounds for the connectivity improvement of partially mobile networks and then present superior empirical results under a variety of different scenarios in which the mobile nodes in our ad-hoc network are embedded with adaptive routing policies and learned movement policies.

Acknowledgements

This work was funded in part by a Graduate Research Fellowship from the National Science Foundation.

I. INTRODUCTION

Mobile ad-hoc networking is emerging as an important research field with a number of increasingly relevant real-world applications, ranging from sensor networks to peer-to-peer wireless computing. Researchers in AI and machine learning have not yet contributed to this growing field. It promises to be a rich and interesting domain for studying the application of learning techniques. It also has direct applications back to AI, for example in the design of communication channels for groups of robots. In this paper, we introduce mobilized ad-hoc networks as a multi-agent learning domain and discuss some motivations for this study. We apply reinforcement learning techniques to two distinct problems within this domain: packet routing and node movement. Using relatively straightforward adaptations of these methods, we are able to demonstrate good empirical results.

Mobile ad-hoc networks have not traditionally been considered a multi-agent learning domain partly because most research in this area has assumed that we have no control over the node movements, limiting research to the design of routing algorithms. Each node is assumed to be attached to some user or object that is moving with its own purpose, and routing algorithms are thus designed to work well under a variety of different assumptions about node mobility patterns.

However, there are many applications in which we can imagine that some of the nodes would have control over their own movements. For example, mobile robots might be deployed in search-and-rescue or military reconnaissance operations requiring ad-hoc communication. In such cases it may be necessary for some nodes to adjust their physical positions in order to maintain network connectivity. In these situations, what we will term *mobilized* ad-hoc networks becomes an extremely relevant multi-agent learning domain. It is interesting both in the variety of learning issues involved and in its practical relevance to real-world systems and technology.

There are several advantages gained by allowing nodes to control their own movement. Stationary or randomly moving nodes may not form an optimally connected network or may not be connected at all. By allowing nodes to control their own movements, we will show that we can achieve better performance for the ad-hoc network. One might view these controllable nodes as “support nodes” whose role is to maintain certain network connectivity properties. As the number of support nodes increases, the network performance also increases. Given better movement and routing algorithms, we can

achieve significant additional performance gains.

It is important to note that there are two levels at which learning can be applied: (1) packet routing and (2) node movement. We will discuss these topics in separate sections in this paper. Packet routing concerns the forwarding decisions each node must make when it receives packets destined for some other node. Node movement concerns the actual movement decisions each node can make in order to optimize the connectivity of the ad-hoc network. Even though we will use reinforcement learning techniques to tackle both these problems, they must be approached with different mindsets. For the routing problem, we focus on the issue of online adaptivity. Learning is advantageous because it allows the nodes to quickly react to changes in network configuration and conditions. Adaptive distributed routing algorithms are particularly important in ad-hoc networks, since there is no centrally administered addressing and routing system. Moreover, network configuration and conditions are by definition expected to change frequently.

On the other hand, the node movement problem is best handled off-line. Learning a good movement policy requires a long training phase, which would be undesirable if done on-line. At execution time, we should simply be running our pre-learned optimal policy. Moreover, this movement policy should encode optimal action selections given different observations about the network state; the overall policy does not change due to changing network configuration or conditions and thus does not need to adapt online. We treat the problem as a large partially-observable Markov decision process (POMDP) where the agent nodes only have access to local observations about the network state. This partial observability is inherent to both the routing and movement portions of the ad-hoc networking problem, since there is no central network administrator. Nodes can only observe the local state around them; they do not have access to the global network topology or communication patterns. Even with this limited knowledge, learning is useful because it would otherwise be difficult for a human designer to create an optimized movement policy for each network scenario.

II. RELATED WORK

We draw inspiration for this work from two different fields: networking and reinforcement learning. In the networking literature, some work on the effect of node mobility in ad-hoc networks has been done for applications in which movement and topology changes are on the time-scale of packet delivery. Nodes are then able

to act as mobile relays physically transporting packets from one location to another. Grossglauser and Tse [1] analyze a strategy in which source nodes send packets to as many different nodes as possible, which store the packets and hand them off whenever they get close to the intended destination nodes. Li and Rus [2] consider a scenario in which mobile hosts make deviations from predetermined trajectories to transmit messages in disconnected networks. Chatzigiannakis et al [3] consider the case where a subset of mobile nodes are constrained to move to support the needs of the protocol, and act as a mobile pool for message delivery.

Our work is in a different setting, in which topology changes are on a much longer time scale than packet delivery delay constraints. Nodes move in order to form and maintain connected routes, rather than to physically deliver packets in a disconnected network. Routing is thus an important aspect of our algorithms that influences and is informed by movement decisions. Related work on routing in mobile ad-hoc networks where no control over node movements is assumed can be found in [4] and [5], while the issue of preserving a connected topology with only some proportion of nodes awake at any one time is addressed in [6].

From the reinforcement learning community, there has been some interest in applying learning techniques to improve network performance. Such adaptive algorithms may be better able to perform well under widely varying conditions. Boyan and Littman [7] applied reinforcement learning techniques to the problem of routing in static networks. They showed that a simple adaptive algorithm based on the Q-learning algorithm [8] can out-perform a shortest-paths algorithm under changing load and connectivity conditions. Peshkin [9] used policy search rather than Q-learning on the same problem, which allowed the system to search a richer policy space. By using stochastic routing policies, the system is able to manage high loads by finding multiple possible source-destination paths. We extend these ideas to the case of mobile networks, where node connectivity is constantly changing.

Moreover, since we assume control over the nodes' movements, we can also influence these connectivity changes by learning a good control mechanism for the node movements. Several papers mentioned above propose various methods for controlling node and packet movement under specific assumptions. In the general setting, we wish to select optimal actions at each time step to maximize the long-term system performance. This type of problem lends itself to reinforcement learning

techniques [10], where the goal of the learner is to maximize long-term reward by learning an optimal behavior policy through simulation. Stone and Sutton [11] and Bowling and Veloso [12] studied methods for scaling up reinforcement learning techniques to complex domains such as robotic soccer.

III. PROBLEM OVERVIEW

Our mobilized ad-hoc network consists of one or more source nodes, one or more receiver nodes, and a number of other wireless nodes within a constrained area. All nodes are independently initialized according to a uniform distribution over this area. The sources generate packets at a constant rate and move according to a random way-point model. The aim of all nodes other than the sources and receivers is to transmit the maximum possible number of packets from the sources to the receivers. Some of these nodes can move so as to aid transmission, while the rest are stationary.

Performance is measured by the proportion of packets successfully transmitted to the receivers. When inter-node link capacities are limited and buffer sizes are finite, packet drops may occur due to buffer overflow, thereby decreasing network performance. When inter-node link capacities are sufficiently large compared to the source packet generation rates, an equivalent performance metric is the average proportion of sources connected to the receivers over time.

The packet transmission success probability achievable in an ad-hoc network depends heavily on various parameters including the transmission range r , the number of network nodes n , and the proportion of mobile nodes m . Alternatively, we may consider what values of these parameters are sufficient to achieve a desired success probability. Our following result illustrates these relationships and gives some sense of the potential benefits of having controllable mobile nodes in an ad-hoc network. It builds upon a theorem by Gupta and Kumar [13] that sets a bound for the minimum (or critical) range necessary for n nodes in a given area to be fully connected. They state that for the case of a disk of area A and n approaching infinity, the critical range for the network is $r_n = \sqrt{\frac{A(\log n + \gamma_n)}{\pi n}}$, where γ_n is a function of n that grows arbitrarily slowly to infinity as $n \rightarrow \infty$. We extend this result to the case of partially mobile networks.

Theorem 1: Let r_n be the minimum, or critical, range needed for n nodes independently and uniformly distributed on a given area to form a fully connected network. If a proportion $m = \frac{k-1}{k}$, k an integer, of the

nodes are mobile, then a transmission range $r = \frac{r_n}{\sqrt{k}}$ is sufficient to make it possible for the mobile nodes to move to form a fully connected network. If the range r is fixed at r_n , with $m = \frac{k-1}{k}$, k an integer, then a total number $\frac{n}{k}$ of nodes suffices for full connectivity.

Proof: If a proportion $m = \frac{k-1}{k}$ of the nodes in a n node network being mobile, these mn nodes can essentially move around to increase the effective transmission ranges of the stationary nodes. For each fixed node, we have $k-1$ mobile nodes. Thus, each stationary node is able to form one link to neighboring node over a distance of kr_n using a set of $k-1$ mobile nodes as relays. Since a fully connected network of n nodes requires only $n-1$ links, each link can be allocated a set of $\frac{(k-1)n}{k(k-1)} > k-1$ mobile relay nodes, allowing each stationary node to transmit a full distance of kr_n rather than r_n . Since $n' = \frac{n}{k}$, from Gupta and Kumar's theorem we know that if all the nodes were stationary, then the critical range for n' nodes would be $r_{n'} = r_{\frac{n}{k}} = \sqrt{\frac{A \log n' + \gamma_{n'}}{\pi n'}} < \sqrt{\frac{A \log n + \gamma_n}{\pi \frac{n}{k}}} = r_n \sqrt{k}$. However, since $(k-1)n$ nodes are now mobile, making the transmission range k times the original range, we only need a range of $r = \frac{r_n}{k}$ for each node; thus a sufficient range for fully connectivity of the partially mobile network is $r = \frac{r_n}{\sqrt{k}}$.

Using the same reasoning, it is easy to show the second part of the theorem. Let us fix the range of the n nodes in the network to be $r_n = \sqrt{\frac{A(\log n + \gamma_n)}{\pi n}}$. We need to show that with $n' = \frac{n}{k}$ total nodes, of which mn' are mobile, each with range r_n , we can form a fully connected network. In this case, we only have $n'' = \frac{n}{k^2}$ stationary nodes. Similar to the above, we know that the critical range for a network with n'' nodes is $\sqrt{\frac{A(\log \frac{n}{k^2} + \gamma_{n/k^2})}{\pi \frac{n}{k^2}}} < r_n k$. Since mn' nodes are mobile, $m = \frac{k-1}{k}$, we actually only need a range of $\frac{r_n k}{k} = r_n$. ■

This result shows that as the proportion of mobile nodes increases, the transmission range r needed for full connectivity decreases for given n (or the minimum value of n required given r decreases). These are loose upper bounds since they allow for up to all links of the fully connected network being greater than $\frac{k-1}{k}$ of the maximum link length.

For more realistic settings with smaller fixed n , we can obtain corresponding empirical results for the minimum range r necessary for achieving nearly full connectivity ($> 95\%$). We ran some trials with n nodes, $25 \leq n \leq 400$, independently initialized according to the

uniform distribution on a unit square. We find that the approximation $r'_n = \sqrt{\frac{A \log n}{\pi n}} \approx r_n$ gives a good estimate of the critical range to within a small decreasing factor. We also calculated the mean proportion h of links that were greater than half the critical range.

n	critical range	mean proportion, h	std. dev.
25	$2.0r'_n$	0.2735	0.1233
50	$1.7r'_n$	0.3265	0.0785
100	$1.7r'_n$	0.2353	0.0647
200	$1.6r'_n$	0.1992	0.0390
400	$1.6r'_n$	0.1588	0.0258

Noting that a substantial proportion of the links are less than half the critical range leads to the conclusion that the bounds given above are quite loose, i.e. the necessary values for each parameter r , n and m , given fixed values for the other two parameters are lower than the sufficient values given by Theorem 1. As such, we would expect potentially better results under conditions of global knowledge and rapid optimal deployment of mobile nodes.

In most realistic settings, nodes only possess local knowledge and movement speed is limited. Thus it becomes harder to reach the potential optimal network performance. Nevertheless, we can develop learning algorithms that perform well given these constraints. In the remainder of the paper, we develop such learning algorithms for movement and routing, and compare their performance against both non-learning and centralized algorithms, where performance is measured by the proportion of packets successfully transmitted under various network scenarios.

IV. Q-ROUTING

To optimize the performance of the ad-hoc network, we need to design good control algorithms for the support nodes. The nodes will need to adapt to changing conditions and communication patterns using intelligent routing and movement algorithms. We focus on the routing problem here and give a brief overview of the movement problem, for which results will be presented in greater detail in a forthcoming paper.

Q-routing [7] is an adaptive packet routing protocol for static networks based on the Q-learning algorithm, which we adapt for use in mobile ad-hoc networks. The algorithm allows a network to continuously adapt to congestion or link failure by choosing routes that require the least delivery time. When a route becomes congested or fails, Q-routing learns to avoid that route and uses

an alternate path. Due to its adaptive nature, we might expect that Q-routing would also work well in the mobile ad-hoc setting.

Q-routing is a direct application of Watkins' Q-learning [8] to the packet routing problem. Q-routing is a distributed routing scheme where each node in the network runs its own copy of the Q-routing algorithm. A node x faces the task of choosing the next hop for a packet destined for some receiver node d . Using Q-routing, it learns the expected delivery times to d for each possible next hop y , where each possible next hop y is a neighbor node connected to x by a network link. Formally, Q-routing keeps Q-tables Q^x for each node x and updates these tables at each time period t as follows:

$$Q_t^x(d, y) = (1 - \alpha)Q_{t-1}^x(d, y) + \alpha(b_t^x + \min_z Q_{t-1}^y(d, z)),$$

where $0 < \alpha < 1$ is parameter that controls the learning rate, and b_t is the time the current packet spent on the buffer or queue at node x before being sent off at time period t .

Q-learning estimates the *value* or *cost*, V , associated with each state d , with $V = \min_z Q^x(d, z)$. In our case, the value of a state is the estimated time for delivery of a packet from the current node x to destination d via node z . Once the nodes have learned the values associated with each state-action pair, they simply execute a greedy policy to behave optimally. When a node receives a packet for destination d , it sends the packet to the neighbor y with the lowest estimated delivery time $Q^x(d, y)$. In experimental trials, Boyan and Littman found that fast learning rates (around 0.5) worked well since it is beneficial for the network to adapt quickly.

Adapting Q-routing to the mobile ad-hoc network routing domain is fairly straightforward. Neighbor nodes are defined as the nodes within transmission range. The main difference is that the neighbor nodes y may appear and disappear quite frequently due to node mobility. When a node y moves out of range, we set the estimated delivery time to d via y to ∞ , i.e. $Q^x(d, y) = \infty$. When a node y moves into range, we optimistically set the estimated time to 0, i.e. $Q^x(d, y) = 0$. This optimistic bias encourages exploration. That is, node x will always try sending packets via a node y that has just come into range. If this action results in a high estimated delivery time, then node x will quickly revert to its original behavior since $Q^x(d, y)$ will quickly be updated to its true value. On the other hand, if this action results in a good delivery time, then node x will continue to send packets via node y .




Fig. 1. A comparison of directional routing vs Q-routing in a network with 10 sources, 15 mobile agents, and one receiver. Simulations were run over 20 initialization positions and 5 source movement scenarios for each different initialization. For each buffer size, averages over all of these trials are depicted, with error bars denoting one standard deviation.

V. EMPIRICAL RESULTS

Our empirical results give an indication of the power of utilizing adaptive learning techniques in designing movement and routing algorithms for mobilized ad-hoc networks. The setup is described in Section III. There are source, receiver, and support nodes in a square grid, usually 30x30 in size. Each node has a transmission range of 6. The support nodes may either be fixed stationary nodes or mobilized agent nodes. There are two time scales: one for transmission and one for movement. During each movement time step, the node can choose one of its movement actions and perform 10 packet transmissions. Each packet transmission is the result of a Q-routing decision and update. Sources generate packets every two transmission time steps, and the number of packets received by a node in any time period is only limited by the node's buffer size.

A. Q-routing vs directional routing

To evaluate the performance of Q-routing, we implement a global knowledge routing policy that is given information about the receiver location. This is done for comparison purposes only; in reality nodes generally do not have access to such information. With this information, nodes can route each packet towards the correct destination by forwarding the packet to a neighboring node that is closest to being in the direction of the receiver. We call this our *directional routing* policy. Specifically, our implementation forwards a packet to the




Fig. 2. Using the directional routing policy, packets often become congested on the trunk of the network tree (Top). Using Q-routing on the same experimental setup (note that the source and receiver nodes are in the same position as the both figures), the mobile nodes in the ad-hoc network spread out to distribute packet load (Bottom). Sources are shown as squares, mobile nodes are circles, and the receiver is an encircled square. Both figures show the simulator after 10000 periods, using the same initialization and movement files.

neighbor that is closest to being in the direction of the receiver, up to a maximum deviation of 90 degrees. If no such neighbor exists, then the packet is dropped.

We compared Q-routing with directional routing under a variety of different scenarios. In almost all cases, Q-routing performs better than directional routing. Especially when the nodes are limited by small buffer sizes, Q-routing performs significantly better. Results for a typical set of network scenarios are shown in Figure 1. This is due to the fact that Q-routing will create alternate paths to the receiver as soon as a path becomes congested. Thus, packets will be less likely to be dropped due to buffer overflow caused by path congestion or limited buffer sizes since alternate paths will be constructed. In directional routing, on the other

hand, often certain paths will become overloaded with traffic, causing significant packet drop.

Q-routing outperforms directional routing even in cases where buffer size is not a direct constraint, such as the case shown in Figure 1 where the buffer size is 10. This illustrates the underlying problem of network capacity. Since the total source packet generation rate exceeds the transmission rate of any one inter-node link, directional routing may still run into trouble with bottleneck links regardless of buffer size. The network capacity achieved using Q-routing is larger since more alternate paths are created around such bottlenecks. Furthermore, directional routing is unable to find circuitous paths from sources to the receiver. Since it only routes packets to neighbors that are in the direction of the receiver, any path that requires a packet to be forwarded away from the receiver for one or more hops will never be found.

These comparisons are done using a fixed movement policy we will call our *centroidal movement* policy. Under this policy, a node that is holding a connection will attempt to move to the centroid of its connected neighbors, which increases the likelihood of preserving these connections over time. If it is not holding a connection, then it simply moves about randomly searching for a new connection. Thus, the next hops determined by the routing policy strongly influence the direction of movement, since the next hops determine the node's connections to its neighbors.

When a random movement policy is used instead of the centroidal policy, Q-routing exhibits inferior performance relative to directional routing. One example is given in Figure 3, which shows the evolution of average system performance over time in a typical scenario. The table below gives averages over 100 different scenarios:

Movement policy	Routing policy	Average performance
Centroidal	Q-routing	.924
Centroidal	Directional	.896
Random	Q-routing	.498
Random	Directional	.519

This phenomenon is due to the fact that Q-routing influences the centroidal movement policy in a positive manner, whereas it is unable to influence a random movement policy. In some sense, Q-routing with centroidal movement is able to find circuitous source-destination paths and rope them in using the centroidal movement.

Due to this type of influence, Q-routing and directional routing result in very different configurations for




Fig. 3. This graph shows a running average of successful transmission rates for a sample network scenario under four cases: Q-routing with centroidal movement, directional routing with centroidal movement, directional routing with random movement, and Q-routing with random movement.

the network when coupled with centroidal movement. Directional routing tends to form a network backbone, which usually comprises the most direct route to the receiver for a large portion of the source nodes. Other sources send packets towards this backbone, resulting in a tree-like network configuration, as shown in Figure 2. Q-routing, on the other hand, always seeks the shortest path towards the receiver, even when buffer sizes are not a constraint. This results in a fan-shaped network configuration, also shown in Figure 2, where each source has its own shortest path to the receiver as long as there are a sufficient number of mobile nodes to create these paths. From this observation, we can begin to see that there is an interplay between the choice of routing protocol and the movement policy of the mobile nodes.

This leads to a subtle but telling explanation for the improved performance of Q-routing over directional routing. In Q-routing, the mobile agent nodes tend to become more dispersed, since no network backbone is created. Thus, as source nodes move about, the Q-routing ad-hoc network is more likely to be able to remain connected without drastic reconfigurations. This interplay between routing and movement forces us to carefully consider the movement policy we choose to pair with our selected routing policy.

VI. LEARNING TO MOVE

The problem of learning a good movement policy is much more difficult. We wish to again apply reinforcement learning techniques to this problem. First of all, the problem of partial observability is much more

pronounced than in the packet routing problem. For example, when the network is in a disconnected state, information about the global network topology is impossible to collect but would be important for determining movements that would optimize network connectivity. Moreover, the local observation space could still be quite large, depending on the observed variables we choose to encode. Secondly, the choice of action space is unclear. At the most basic level, the agents could move in any direction at a given speed. We will limit the action choices by designing more complex actions that incorporate domain knowledge. This section briefly outlines our application of Q-learning to learn a reasonable movement policy despite the fact that Q-learning generally fails in POMDPs.

We proceed by using the observation space as our state space. This can potentially lead to problems of aliasing, but we choose our observed variables carefully in order to avoid this pitfall. Since the nodes communicate using a shared wireless medium, a node can “sniff” packets sent by neighbors to destinations other than itself. Thus, a node can detect the presence of neighbor nodes and their connections, even if it is not involved in any of these connections itself. Moreover, since the receiver nodes send back acknowledgement packets along these connections, our agent node can also collect statistics about these source-destination paths by sniffing these acknowledgement packets. Each agent node’s observation space thus includes the identities of its neighbors, the number of connections it is currently holding, the number of nearby connections, and the maximum and minimum average hop lengths of these source-destination paths.

In order to constrain the learning process, we incorporate some domain knowledge into our design of an appropriate action space. For example, there is little need to train the nodes to avoid obstacles along its desired path of movement. We can pre-program the nodes with the necessary algorithm to do this. Learning is focused on higher-level action selection that is difficult to pre-program optimally. A subset of our action space is given in the table below. Many of these actions could take multiple time periods to complete. We get around this problem by allowing the agents (nodes) to choose to either continue or change an action during each time period.

Action	Description
stay	Stay put; don't change position.
direction	Head in a particular direction.
plug	Searches for the sparsest path and attempts to fill in the largest gap along that path.
leave	Leaves a path and becomes available for other actions.
circle	Circles around a node in search of more connections. Attempts to retain connection to the source around which it is circling.
lead	Leads other mobile agents in search of more connections. Attempts to stay within range of its followers.
follow	Identifies and follows a leader node, while maintaining connection to previous hop.
center	Moves to the centroid of the neighbors to which it is connected.
explore	Random exploration.

Finally, the reward given to the nodes during each time period corresponds to the percentage of successful transmissions during that time period, which is available since we are conducting this training off-line.

A. Movement policy comparisons

We evaluate the performance of our learning algorithm against the centroidal movement policy given in Section V, a hand-coded policy that uses the same observation space as the learning algorithm, and a global knowledge central controller. Under simulation, we give the central controller access to all the node, source, and receiver positions, which would usually be unavailable to the agent nodes. Since our learning algorithm only has access to local knowledge, the central controller's performance should approximate an upper bound for the learning algorithm's performance. Moreover, this performance bound may fluctuate over time as network conditions change depending on source movement scenarios. The central controller also gives us a mechanism to compare our theoretical performance bounds from Section III with actual empirical results.

The central controller is designed to approximate an optimal movement policy given global knowledge about network conditions. It begins by identifying connected components among the stationary nodes. If a packet is received by any node of a connected component, then all nodes of that component can also receive the packet. However, if a packet's destination is in a different




Fig. 4. Graph showing the average performance of various movement policies over time in a typical scenario. The learning policy is shown during its training phase.

component, then the controller finds the shortest path to the destination component and recruits mobile nodes to help construct the necessary links needed to deliver the packet.

Figure 4 gives the average performance of a sample network scenario over time using each of these movement policies. As we can see, the learning algorithm does eventually learn a policy that behaves fairly well, but it never achieves the performance of the global knowledge controller. This is expected since the learner does not have access to the global network topology. On average, the learned policies perform slightly better than the hand-coded policy over large sets of network scenarios. However, in certain scenarios, it never learns to perform as well, possibly due to aliasing problems.

VII. CONCLUSION AND FUTURE WORK

This paper presents a rich new domain for multiagent reinforcement learning and establishes several first results in this area. There is much more work to be done to create a more robust movement learning algorithm that deals directly with the partial observability of the domain. Policy search and Monte Carlo methods are two techniques that are currently being examined. A well designed algorithm might also be able to exploit the interplay between routing and movement.

REFERENCES

- [1] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In *INFOCOM*, 2001.
- [2] Q. Li and D. Rus. Sending messages to mobile users in disconnected ad-hoc networks. In *MOBICOM*, 2000.

- [3] I. Chatzigiannakis, S. Nikolettseas, N. Paspallis, P. Spirakis, and C. Zaroliagis. An experimental study of basic communication protocols in ad-hoc mobile networks. In *5th Workshop on Algorithmic Engineering*, 2001.
- [4] C. Perkins and E. Royer. Ad-hoc on-demand distance vector routing. In *MILCOM Panel*, 1997.
- [5] D. Johnson and D. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, volume 353. 1996.
- [6] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: An energy-efficient coordination algorithm for topology maintenance. In *SIGMOBILE*, 2001.
- [7] J. Boyan and M. L. Littman. Packet routing in dynamically changing networks: A reinforcement learning approach. In *Advances in NIPS*, 1994.
- [8] C. J. Watkins. Learning with delayed rewards. *Ph.D. Thesis, University of Cambridge*, 1989.
- [9] L. Peshkin. Reinforcement learning by policy search. *Ph.D. Thesis, MIT*, 2002.
- [10] L. P. Kaelbling, M. L. Littman, and A. P. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- [11] P. Stone and R. S. Sutton. Scaling reinforcement learning toward RoboCup soccer. In *ICML*, 2001.
- [12] M. Bowling and M. Veloso. Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136:215–250, 2002.
- [13] P. Gupta and P. R. Kumar. Capacity of wireless networks. In *Stochastic Analysis, Control, Optimization, and Applications*. Birkhauser, 1998.