

Stochastic Routing in Ad Hoc Networks

Christopher Lott

Qualcomm, Inc., San Diego, CA 92121 (858)651-4393 clott@qualcomm.com

Demosthenis Teneketzis

Department of EECS, University of Michigan, Ann Arbor, MI 48109 (734)763-0598 teneketzis@eeecs.umich.edu

Abstract—We investigate a network routing problem where a probabilistic local broadcast transmission model is used to determine routing. We discuss this model's key features, and note that the local broadcast transmission model can be viewed as soft handoff for an ad hoc network. We present results showing that an index policy is optimal for the routing problem. We extend the network model to allow for control of transmission type, and prove that the index nature of the optimal routing policy remains unchanged. We present three distributed algorithms which compute an optimal routing policy, discuss their convergence properties, and demonstrate their performance through simulation.

I. INTRODUCTION

The term *ad hoc* is applied to networks in which there is no central network controller; each node can itself act as a store-and-forward router, and in which the connection topology is time-varying (e.g. see [11]). Such a network is in contrast to, for example, a cellular network where each cell has a central base station through which all cell data is transmitted. An ad hoc network contains a number of packets, each of which is destined for some set of destination nodes. The general routing problem is to define a policy which, given the trajectory histories of all the packets, chooses which nodes should next transmit which packet. It is usually desirable that this policy be implemented in a distributed form, so transmission decisions can be decided locally without knowledge of other parts of the network.

Many approaches are possible for network routing optimization. A typical one is maximizing the overall network throughput. But in many cases other considerations are at least as important. In the case of a wireless network the energy source is locally stored in a battery at each node, and a major design goal is often to achieve satisfactory communication while using up as little energy as possible. This low-energy requirement takes an extreme form in *sensor networks*, where large arrays of mobile information-gathering devices must communicate under severe energy limitations [2].

This paper explores some design issues in network routing algorithms for ad hoc networks, and provides a novel system model which allows for optimal design in a number of realistic situations. It is organized as follows. In the remainder of Section I we present briefly a discussion of the literature available on routing in ad hoc wireless networks, and a summary of the contributions of this paper. In Section II we define and solve a stochastic routing problem. In Section III we

extend this routing problem to include choice of transmission type. In Section IV through Section VIII, we present three distributed algorithms which compute an optimal routing policy, and discuss their convergence properties and performance. Section IX provides conclusions and ideas for future work.

A. Ad Hoc Network Routing Literature

There has been much recent research activity on routing for ad hoc networks. This work can perhaps most conveniently be categorized as either *on-demand* or *route maintenance* (also sometimes referred to as *reactive* and *proactive* [23]).

1) *On-Demand Protocols*: On-demand protocols are the algorithms which only update routing tables when new packets arrive. On-demand protocols usually store old routing information and use it as long as packets still get to their destination. When a given path fails, there is some mechanism for probing the network to find a new path, which often involves flooding the network. Link detection used to decide path failure generally uses information from the Medium Access Control (MAC) at the link layer. Usually the path found is not optimal in any sense. The value of an on-demand protocol is that the communication requirement to set up a route is minimized. The idea is that in general it would cost more to find a better path than the improved route would reduce cost. Hence, these protocols can be valuable when network dynamics are rapid compared to packet transmission frequency.

A number of on-demand routing protocols have been proposed for ad hoc networks. The Temporally-Ordered Routing Algorithm (TORA) [6], [22] is based on work originally in [8]. TORA tries to quickly establish routes on demand using a path reversal method which is guaranteed to find some route, with no criteria for the quality of the route. Dynamic Source Routing (DSR) [13], [19] uses source routing, where each packet header contains the entire route for the packet. Existing routes are used without update until one fails. When one fails, a Route Discovery is executed, during which Route Request packets flood the network until any route is found. The Ad Hoc On-Demand Distance Vector (AODV) protocol [25] is a kind of hybrid of on-demand and route maintenance protocols. The link status between a node and its neighbors is maintained through periodic beacon signals between nodes. But routes are only updated when a packet routing fails, and then this link status information is used. The route chosen by AODV is the one with the minimum number of hops. The Zone Routing

Protocol (ZRP) [23] is another hybrid which allows trading off of route update and route maintenance overheads. The network nodes are partitioned into zones. As in route maintenance, each node maintains a continuous knowledge of the topology of all the nodes within its own zone. However, routes to nodes outside a node's zone are determined on-demand in a Route Discovery process. Independent Zone Routing (IZR) is an update to ZRP which allows zone size to vary per node [31].

2) *Route Maintenance Protocols*: As the name implies, route maintenance protocols periodically expend the energy and bandwidth necessary to maintain updated routing tables under all traffic conditions. Given their wide acceptance, the benchmarks for route maintenance protocols are often the TCP/IP standards RIP (Routing Information Protocol) [18] and OSPF (Open Shortest Path First) [21]. However the standard Internet implementations of these algorithms are not well-suited to the ad hoc network environment. In highly mobile environments, communication overhead to implement a full link-state algorithm is prohibitively costly. Alternatively, much research has been performed to adapt distance vector type algorithms to the mobile environment. In these approaches, it is common to use a Distributed Bellman-Ford (DBF) type algorithm [4] to update the routes and cost estimates. A common theme in the research, then, is to address the weaknesses of this type of algorithm in a highly mobile and uncertain environment.

An early effort along these lines is [20], which endeavors to provide loop-free routing in a quasi-static environment. In the Destination-Sequenced Distance Vector (DSDV) protocol [24], routes are tagged with sequence numbers. This allows the most up-to-date route information to be included, as well as providing a quick means for propagating link breakage information. In this way, loop-free routing is achieved. In Least Resistance Routing (LRR) [26], [27], the link quality between nodes and the current buffer length of the receiving node are used to define a link resistance value. A DBF type algorithm is then used to find a least resistance path for such a network. The Loop-free Path-finding algorithm (LPA) [9] eliminates the counting-to-infinity and packet looping problems with the DBF algorithm by using predecessor information and loop-detection at each router.

A minimum energy approach to routing in ad hoc networks is taken in [28]. The most striking difference to our work is in the channel model. In [28] communication is peer-to-peer, and the stochastic nature of the channel is modeled as a threshold value for transmission power that ensures reliable communication. This threshold value is computed from a signal transmission model requiring knowledge of node positions (via global positioning system (GPS) fixes). Once energy requirements to transmit to each node are determined, a DPF algorithm is run to compute minimum energy routes.

An interesting approach to the multicast problem is developed in [33] and [34]. A deterministic local broadcast model for the link, with a power threshold determining the availability of a link, is used to set up a problem to find the minimum energy to span a given set of destination nodes. An optimal solution to this problem is computationally infeasible, but the satisfactory performance of heuristic algorithms is studied

using simulation. This circuit-switched multicast problem with deterministic local broadcast model is quite different from the problems we define in this paper.

A detailed comparison via simulation of a few different routing protocols for ad hoc networks is described in [5]. To truly understand the tradeoffs involved among the wide variety of available algorithms, more such work is necessary. Even better would be progress in our conceptual understanding, leading to better methods for critiquing the different approaches.

For a useful guide to the large and growing literature on ad hoc networks, see [11] and the references therein.

B. Contribution

In this work we present what is, to the best of our knowledge, the first network routing protocol which uses a probabilistic local broadcast model for transmission. This model (i) provides an intimate coupling between lower network layers through careful modeling of the key channel characteristics, and (ii) allows for routing decisions to be made based on immediate feedback from each transmission. In most network models in the literature, nodes are connected by links, and the routing algorithm dictates which set of links each packet is to traverse. This is true even when through local broadcast a packet can traverse multiple links in parallel. Typically at each point in time, through an estimate of its SNR, a link is designated as deterministically up or down, with SNR sometimes used to determine the sustainable link rate. The stochastic transmission nature is subsumed in a link cost or delay, or in a success probability for the entire link layer transmission protocol. In contrast, our protocol detects and reacts to the stochastic result of each node's local broadcast transmission, and constructs an optimal route based on this immediate feedback. In consequence, the actual route taken by a packet depends on random system events, may vary among packets heading to a common destination, and is fully determined only via actual transmission. This reacting to individual transmission results can be interpreted as the ad hoc network analog of soft handoff in a cellular system, as discussed further below. We show that such a protocol is optimal for the stochastic local broadcast model. We also show the protocol is well suited to a distributed implementation, as it is based only on information local to a node.

We extend the model to allow for control of transmission type at each node, and show how the fundamental nature of this optimal protocol doesn't change. Varying transmission type to optimize routing decisions couples physical layer considerations with the network layer routing function. For example, in the particular case of power control, an optimal protocol effectively resolves the tradeoffs between fewer long hops vs. more short hops.

We present three distinct algorithms which compute an optimal index policy for our stochastic routing problem in an asynchronous distributed fashion, so that at each computation only information local to each node is used. The algorithms are all similar in the asynchronous nature of their value updates and transmissions, with the fundamental difference being how the update value is computed. The algorithm presented in

Section V uses an update motivated by the stochastic Dijkstra algorithm we introduce in Section II. The algorithm presented in Section VI implements Distributed Dynamic Programming (DDP), as defined in [3], and this leads to an alternate update computation. Finally, the algorithm presented in Section VII, which we call the rank algorithm, is a third way to implement the update. For each algorithm we prove the same fundamental convergence result: 1) Asymptotic convergence of the value function, 2) Finite time convergence to an optimal policy, and 3) Finite time convergence to an optimal index policy under the condition that all nodes have distinct value functions. We end by comparing the performance of these three algorithms with other approaches in a realistic simulation environment.

There is an interesting precedent for the stochastic local broadcast model used in this paper. In CDMA cellular systems [32], soft handoff is used to improve network efficiency. On the reverse link, even though the mobile is under transmit power control, random channel variation makes it uncertain which base station will successfully decode each mobile packet. Efficient operation is achieved by having each base station attempt to decode each packet, and the packet is successful if any of them succeed. In effect, each base station is a destination node and the random result of each mobile local broadcast is used to decide on packet success. Our stochastic routing problem can be viewed as a generalization of soft handoff to the case of an ad hoc network, where intermediate store-and-forward nodes may also receive each packet.

In CDMA cellular system reverse links [7][16], a mobile's transmit power is controlled to achieve a fixed packet error rate (PER). Implicitly this relies on the fact that the recent interference level seen by mobile transmissions is a good estimator of future interference. No further attempt is made to coordinate mobile transmissions across the network (interference cancellation is a separate issue). In our approach here, rather than controlling mobile transmit power to achieve a fixed PER, we assume a given set of transmit power levels, and estimate the resulting packet success probabilities to neighbors. By taking this approach, we also implicitly rely on some persistence in interference, and hence transmission success, statistics. And in deference to the inherently distributed nature of an ad hoc network, we also do not attempt any further coordination of mobile transmissions, so that each mobile views other mobile transmissions solely as random sources of interference.

Stochastic routing also can be viewed as a means to achieve dynamic network load balancing. In effect, each mobile transmission is a parallel attempt across neighbor channels, and packet success depends on the concurrent loading. Packets thus are naturally routed toward less loaded portions of the network, with minimal retransmissions. For longer-term loading imbalances, the transmission probability estimates themselves reflect interference-limited regions of the network. These aspects of stochastic routing are important for the network simulation results presented in Section VIII.

II. THE STOCHASTIC ROUTING PROBLEM

A. Notation and Preliminaries

We begin by briefly summarizing notation and definitions for the system model under consideration, which we refer to

as Model **(M)**. In Model **(M)** control is centralized, meaning the controller has access to all information in the network. Also, Model **(M)** is probabilistic, with transitions described by $P(S|i)$.

N is the number of nodes in the network.

$\Omega = \{1, \dots, N\}$, the set of all nodes. So $|\Omega| = N$.

$S \subseteq \Omega$ refers to a state of the system, defined as the set of nodes which have received the packet. S_t refers to the state at time t .

$R : 2^\Omega \rightarrow \mathbb{R}$ is the reward function, and $R_i := R(\{i\})$. Also $R_{max} := \max_{i \in \Omega} R_i$.

π is a Markov policy. We write $\pi(S) = i$ to indicate policy π transmits at node i when in state S .

We write $\pi(S) = r$ to indicate policy π retires and receives reward $R(S)$ when in state S . For convenience we write $\pi(S) = r_i$ as shorthand that policy π retires and receives $R_i, i \in S$. In this case, we say that policy π retires and *receives the reward* of node i .

By $\pi(S) \neq i, r_i$, we mean both $\pi(S) \neq i$ and $\pi(S) \neq r_i$.

By $\pi(S) = \tilde{\pi}(S)$, we mean either $\pi(S) = \tilde{\pi}(S) = i$, or $\pi(S) = \tilde{\pi}(S) = r_i$, for some i .

$V^\pi(S)$ is the expected reward when starting in state S under policy π . We often write V_i^π for $V^\pi(\{i\})$. When π is optimal, we use $V^\pi(\cdot)$ to indicate the optimal value function.

We write $P^i(S'|S)$ to indicate the probability of reaching state S' from state S when choosing i for transmission, $i \in S$. We write $P^i(S|i)$ as shorthand for $P^i(S|\{i\})$.

We define $P_{ij} := \sum_{S: i, j \in S} P^i(S|i)$.

j is called a *neighbor* of i if $P_{ij} > 0$. $\mathcal{N}(i)$ is the set of all neighbors of i , together with i itself. Note that $P_{ij} \neq P_{ji}$ is permitted.

By $\operatorname{argmax}_{i \in S} f(i)$, we mean the *set* of values of i from the finite set S which maximizes $f(i)$.

Definition 2.1 (Increasing Property): Model **(M)** is said to be *increasing* if for any system realization under any policy we have $S_{t_2} \supseteq S_{t_1}, \forall t_1, \forall t_2 > t_1$.

Definition 2.2 (Decoupling Property): Model **(M)** is said to be *decoupled* if

$$P(S_2|S_1) = \sum_{S \subseteq S_3} P(S_2 - S | S_1 - S_3), \quad \forall S_1, S_2 \subseteq \Omega, \quad \forall S_3 \subset S_1.$$

The meaning of this definition is that transmission success to a set of neighbors from a node at a given time is unaffected by which other nodes already have the packet.

Definition 2.3: A function $f : 2^\Omega \rightarrow \mathbb{R}$ is an *index function* on Ω if f satisfies

$$f(S) = \max_{i \in S} f(\{i\}) \quad \forall S \subseteq \Omega \quad (1)$$

We next formulate the centralized version of the stochastic routing problem.

B. Statement of Problem

Problem **(P₁)**

We consider the transmission of a single packet, from a given initial state S_o (i.e. a given set of nodes) to a set of destination states, in a wireless ad hoc network of N nodes described by Model **(M)**. Transmission instances occur at

discrete time points. Each transmission from a given node i incurs a fixed cost $c_i > 0$. According to Model (M): (i) at each transmission instance the transmitting node is chosen by a central controller that always knows the current state of the system (i.e. the set of nodes that have the packet); (ii) node transmissions are local broadcasts, that is, multiple neighbor nodes may all simultaneously receive the packet; (iii) given the node chosen to transmit, the probability that a given set of nodes receives the packet is known and fixed; (iv) The central controller is informed, without any cost, as to which nodes receive the packet. In general, control information flow between the nodes and the controller is considered free in energy and instantaneous in time; and (v) each transmission event is assumed independent of those before and after. We assume Model (M) is increasing and decoupled. A reward function R is specified, where R is an index function. At each instance, the central controller chooses either to terminate the transmission process or to continue transmitting. The objective is to choose: (i) the node to transmit at each transmission instance, and (ii) the instance to terminate the transmission process, to maximize

$$E \left\{ R(S_f) - \sum_{t=1}^{\tau-1} c_{i(t)} \right\} \quad (2)$$

where τ is the time when the transmission process is terminated, S_f is the state at τ , and $i(t)$ is the node chosen by the transmission policy at time t .

C. Analysis of Problem (P₁)

We analyze Problem (P₁) and discuss the character of an optimal policy π . The system of Problem (P₁) is a time-homogeneous Markov chain, hence we are faced with a finite-state Markovian Decision Problem with perfect information. We can thus restrict attention to Markov policies on 2^Ω , and we are guaranteed that such an optimal Markov policy exists (cf. [29] Ch.3 p.51). We seek an optimal Markov policy $\pi : 2^\Omega \rightarrow (1, \dots, N)$ which minimizes (2).

To solve Problem (P₁), we could directly apply stochastic dynamic programming. But since the number of states is 2^N , the complexity of such an approach is at least $O(2^N)$, and generally higher (see [10]). Instead, we use the special structure of this problem to find a better algorithm.

1) *Structure of an Optimal Policy for Problem (P₁):* We begin with some definitions.

Definition 2.4: A Markov policy π is a *priority* policy if there is a strict priority ordering of the nodes s.t. $\forall i \in \Omega$ we have $\pi(S \cup \{i\}) = \pi(\{i\}) = i$ or $r_i, \forall S \subseteq \Omega_i$, where Ω_i is the set of nodes of priority lower than i .

Definition 2.5: For priority policy π , we write $i >_\pi j$ when i has higher priority than j under π .

Definition 2.6: A priority policy π is called an *index* policy if $V^\pi(\cdot)$ is an index function on Ω .

Note that a priority policy need not be an index policy.

Our main goal in this section is to prove that there exists an index policy which is an optimal Markov policy for Problem (P₁). We state this result in the following theorem.

Theorem 2.1 (Index Policy): There is an optimal Markov policy π for Problem (P₁) which is an index policy.

We develop a series of lemmas which are used to prove Theorem 2.1. In the first lemma we show that the definition of an index function is equivalent to requiring two properties on f .

Lemma 2.1: Function f is an index function on Ω if and only if the following two properties (3) and (4) both hold.

$$f(S \cup \{i\}) \geq f(S), \quad \forall S \subset \Omega, i \in \Omega \quad (3)$$

$$f(S) \neq f(\{i\}) \implies f(S) = f(S - \{i\}), \quad \forall S \subseteq \Omega, i \in S \quad (4)$$

Proof: Assume f is an index function on 2^Ω . Then f can be written in the form (1). We have $\forall S \subseteq \Omega, i \in \Omega$

$$f(S \cup \{i\}) = \max_{j \in S \cup \{i\}} f(\{j\}) \geq \max_{j \in S} f(\{j\}) = f(S) \quad (5)$$

and (5) establishes (3). To establish (4), assume we have an $i \in S$ where $f(S) \neq f(\{i\})$. Then

$$f(S) = \max_{j \in S} f(\{j\}) = \max_{j \in S - \{i\}} f(\{j\}) = f(S - \{i\}) \quad (6)$$

Together (5) and (6) establish that if f is an index function, then (3) and (4) both must hold.

Conversely, assume (3) and (4) hold. We proceed by induction on the number of elements in S . When $|S| = 1$, for any S it is clear that $f(\{i\}) = \max_{j \in \{i\}} f(\{j\})$, so the induction base step is established. Now assume $f(S)$ can be written in the form of (1) $\forall S \subset \Omega$ s.t. $|S| = K$. Let $S' \subseteq \Omega$ be s.t. $|S'| = K + 1$. We consider two cases.

Case 1 $f(\{i\}) = f(\{j\}) \forall i, j \in S'$

Assume there is an $i \in S'$ s.t. $f(S') \neq f(\{i\})$. Then by (4) and using the fact that $|S' - \{i\}| = K$ with the inductive hypothesis, we have $f(S') = f(S' - \{i\}) = \max_{j \in S' - \{i\}} f(\{j\}) = f(\{i\})$, which is a contradiction. So we must have $f(S') = f(\{i\}) = \max_{j \in S'} f(\{j\})$, and (1) holds for this case.

Case 2 $\exists i, j \in S'$ s.t. $f(\{i\}) < f(\{j\})$

By (3), $f(S') \geq f(\{j\}) > f(\{i\})$. Then by (4), and using the fact that $|S' - \{i\}| = K$ with the inductive hypothesis, we have

$$f(S') = f(S' - \{i\}) = \max_{j \in S' - \{i\}} f(\{j\}) = \max_{j \in S'} f(\{j\}) \quad (7)$$

and (1) holds for this case.

In both cases f has the form of (1) for S' , and this completes the induction step. So by induction f must be an index function on 2^Ω . ■

Next, we use the decoupling and increasing properties of Problem (P₁) to show that the optimal value function for Problem (P₁) possesses a monotonicity property.

Lemma 2.2: (Monotonicity) In Problem (P₁), let π be an optimal Markov policy, and let $S_1, S_2 \subseteq \Omega$ and $S_2 \subseteq S_1$. Then $V^\pi(S_2) \leq V^\pi(S_1)$.

Proof: Given π and $S_2 \subseteq S_1$, we define a new policy $\hat{\pi}$ acting on state S_1 as follows. Let $i = \pi(S_2)$, and suppose S_4 is the state resulting if at first π were to choose i . At the first step, $\hat{\pi}$ chooses i from S_1 , which is possible since $S_2 \subseteq S_1$. $\hat{\pi}$ learns the result of the transmission, and hence knows the new actual state of the system, which we call S_3 .

Furthermore, since $\hat{\pi}$ knows which nodes receive the packet even if they already have it, it also knows what the new state would be if the previous state were S_2 instead of S_1 . By the decoupling property, this new state is S_4 . This fact together with the increasing property also imply that $S_4 \subseteq S_3$, since $S_2 \subseteq S_1$.

At the next step $\hat{\pi}$ acts on S_3 by choosing the same node as π would use on S_4 ; this is possible because $\hat{\pi}$ knows S_4 , and $S_4 \subseteq S_3$. The process continues in this way until $\hat{\pi}$ retires at the same time at which π would retire. Policy $\hat{\pi}$ knows π 's retirement time because it knows the state π "sees" at each time. Let S_{f1} and S_{f2} be the states at retirement for $\hat{\pi}$ and π , respectively. By the above argument, we know $S_{f2} \subseteq S_{f1}$. At retirement, $\hat{\pi}$ has incurred the same cost as π , since π and $\hat{\pi}$ use the same nodes to transmit. Because R is an index function, by Lemma 2.1 (3) we have

$$V^{\hat{\pi}}(S_1) - V^{\pi}(S_2) = R(S_{f1}) - R(S_{f2}) \geq 0 \quad (8)$$

Because π is optimal and $\hat{\pi}$ is suboptimal, we conclude from (8) that

$$V^{\pi}(S_1) \geq V^{\hat{\pi}}(S_1) \geq V^{\pi}(S_2) \quad (9)$$

This completes the proof. ■

In the next lemma, we use the increasing property of Problem (P₁) to show that a Markov policy which is optimal for all states that are a superset of some $S_1 \subset \Omega$, and that takes an optimal action when in S_1 , is also optimal when in state S_1 .

Lemma 2.3: Let $\tilde{\pi}$ be an optimal Markov policy for Problem (P₁). Suppose we are given S_1 , and let π be a Markov policy which has the following two properties.

$$V^{\pi}(S) = V^{\tilde{\pi}}(S), \quad \forall S \supset S_1 \quad (10)$$

$$\pi(S_1) = \tilde{\pi}(S_1) \quad (11)$$

Then

$$V^{\pi}(S_1) = V^{\tilde{\pi}}(S_1) \quad (12)$$

Proof: If $\pi(S_1) = \tilde{\pi}(S_1) = r_i$, for some $i \in S_1$, then (12) holds.

Suppose $\pi(S_1) = \tilde{\pi}(S_1) \neq r$. We compare π and $\tilde{\pi}$ when both transmit in state S_1 . Let \tilde{S}_2 and S_2 be the state after transmitting when in S_1 for $\tilde{\pi}$ and π , respectively. Due to (11), we have

$$\tilde{S}_2 = S_2 \quad (13)$$

Due to the increasing property, we have

$$\tilde{S}_2 = S_2 \supseteq S_1 \quad (14)$$

By (10), we have

$$V^{\tilde{\pi}}(\tilde{S}_2) = V^{\pi}(S_2), \quad \tilde{S}_2 = S_2 \supset S_1 \quad (15)$$

Equations (13)-(15) mean that $\tilde{\pi}$ and π choose the same node from S_1 for transmission, and either reach the same state $S_2 \supset S_1$, which has the same value function for both policies, or both stay in $S_2 = S_1$, at which point they again both play the same node for transmission. Hence, (12) follows. ■

Next, we construct a Markov policy for Problem (P₁) which has many characteristics necessary for an index policy, and then use the lemmas presented above to show that this policy is optimal. The result of this lemma is instrumental in proving Theorem 2.1.

Lemma 2.4: Let $\tilde{\pi}$ be an optimal Markov policy for Problem (P₁). Then there exists a Markov policy π which has the following three properties.

1) For all $S \subseteq \Omega$ where $|S| \geq 2$,

$$\pi(S) = i \implies \pi(S - \{j\}) = i \quad \forall j \in S, j \neq i \quad (16)$$

$$\pi(S) = r_i \implies \pi(S - \{j\}) = r_i \quad \forall j \in S, j \neq i \quad (17)$$

2) For all $S \subseteq \Omega$ where $|S| \geq 2$, and $\pi(S) = i, r_i$,

$$V^{\pi}(S - \{j\}) = V^{\pi}(S) = V^{\tilde{\pi}}(S) = V^{\tilde{\pi}}(S - \{j\}) \quad \forall j \in S, j \neq i \quad (18)$$

3) π is an optimal Markov policy.

Proof: We define π using the following rules:

$$\pi(\Omega) = \tilde{\pi}(\Omega) \quad (19)$$

$$\pi(S - \{j\}) = \pi(S), \quad \forall S \subseteq \Omega, \forall j : \pi(S) \neq j, r_j \quad (20)$$

$$\pi(S - \{j\}) = \tilde{\pi}(S - \{j\}), \quad \forall S \subseteq \Omega, j : \pi(S) = j, r_j \quad (21)$$

If $|\Omega| = N = 1$, by (19) the lemma is true. Assume $|\Omega| = N \geq 2$.

It follows directly from (19)-(21) that π satisfies (16) and (17).

We prove (18) by backward induction on the cardinality of S . By the optimality of $\tilde{\pi}$, we know that $\tilde{\pi}(\Omega) = r_i$ for some $i \in \argmax_{j \in \Omega} R_j$. By (19) and (20) we have $\pi(\Omega) = \pi(\Omega - \{j\}) = r_i, \forall j \in \Omega, j \neq i$. That is, π acting on both Ω and $\Omega - \{j\}$ immediately retires and receives reward R_i . We also have $V^{\tilde{\pi}}(\Omega) = R_i = V^{\tilde{\pi}}(\Omega - \{j\})$, because $\tilde{\pi}$ is optimal and i is available for retirement in $\Omega - \{j\}$. Hence, when $\tilde{\pi}(\Omega) = r_i$, we have $\forall j \in \Omega, j \neq i$

$$V^{\tilde{\pi}}(\Omega - \{j\}) = V^{\tilde{\pi}}(\Omega) = V^{\pi}(\Omega) = V^{\pi}(\Omega - \{j\}) = R_i \quad (22)$$

Equality (22) proves (18) for π when $S = \Omega$, and the basis for induction is established.

If $N = 2$, the argument of (22) completes the proof of (18). We now assume $N > 2$ and prove the induction step. Assume (18) is true for any state S where $|S| \geq L + 1, 2 \leq L < N$. Consider any state S_1 where $|S_1| = L$.

We first prove that

$$V^{\pi}(S_1) = V^{\tilde{\pi}}(S_1) \quad (23)$$

If there exists $j \in \Omega - S_1$ such that $\pi(S_1 \cup \{j\}) \neq j, r_j$, then by (20) we have $\pi(S_1) = \pi(S_1 \cup \{j\})$. By the induction hypothesis, equation (18) is true for $S = S_1 \cup \{j\}$, because $|S_1 \cup \{j\}| = L + 1$. We thus have

$$V^{\pi}(S_1 \cup \{j\} - \{j\}) = V^{\tilde{\pi}}(S_1 \cup \{j\} - \{j\}) \quad (24)$$

Equation (23) follows from (24).

If no such j exists, then by (21) we have $\pi(S_1) = \tilde{\pi}(S_1)$. Because $\pi(S_1) = \tilde{\pi}(S_1)$, and $V^{\pi}(S) = V^{\tilde{\pi}}(S), \forall S \supset S_1$ (because of the induction hypothesis), the conditions (10) and (11) of Lemma 2.3 are satisfied by π and $\tilde{\pi}$ for S_1 . Hence

$$V^{\pi}(S_1) = V^{\tilde{\pi}}(S_1) \quad (25)$$

We have shown that (23) holds for any S_1 where $|S_1| = L$. We use (23) to show that (18) holds for all S_1 where $|S_1| = L$. For the remainder of the proof, assume that either $\pi(S_1) = i$ or $\pi(S_1) = r_i$, and let $j \in S_1, j \neq i$.

Consider first the case where $\pi(S_1) = r_i$. By (20), $\pi(S_1 - \{j\}) = r_i$, so that

$$V^\pi(S_1 - \{j\}) = V^\pi(S_1) = R_i \quad (26)$$

By Lemma 2.2 and the optimality of $\tilde{\pi}$, we have

$$V^{\tilde{\pi}}(S_1 - \{j\}) \leq V^{\tilde{\pi}}(S_1) \quad (27)$$

By (23),

$$V^{\tilde{\pi}}(S_1) = V^\pi(S_1) = R_i \quad (28)$$

But $i \in S_1 - \{j\}$, and $\tilde{\pi}$ is an optimal policy, so

$$V^{\tilde{\pi}}(S_1 - \{j\}) \geq R_i \quad (29)$$

Relations (27), (28), and (29) together imply that

$$V^{\tilde{\pi}}(S_1 - \{j\}) = V^{\tilde{\pi}}(S_1) = R_i \quad (30)$$

Equations (23), (26), and (30) imply that

$$V^{\tilde{\pi}}(S_1 - \{j\}) = V^{\tilde{\pi}}(S_1) = V^\pi(S_1) = V^\pi(S_1 - \{j\}) = V^{\tilde{\pi}}(S_1 - \{j\}) \quad (31)$$

and the induction step for (18) is proved when $\pi(S_1) = r_i$.

Now consider the case where

$$\pi(S_1) = i \quad (32)$$

We claim that

$$\pi(S) \neq j, r_j \quad \forall S \supseteq S_1 - \{j\} \quad (33)$$

We prove (33) as follows. Let S be any state where $S \supseteq S_1 - \{j\}$. If $j \notin S$, then $\pi(S) \neq j, r_j$, and (33) follows. Assume $j \in S$. Then $S \supseteq S_1$ and $|S| \geq L$. If $|S| = L$, then $S = S_1$ and $\pi(S) = i \neq j, r_j$. Assume $|S| > L$. If $\pi(S) = j$, then by successively removing nodes $k \in \Omega - S_1$ and using (20) we obtain $\pi(S_1) = j$, which contradicts (32). If $\pi(S) = r_j$, then by successively removing nodes $k \in \Omega - S_1$ and using (20), we obtain $\pi(S_1) = r_j$, which contradicts (32). Hence (33) is true in all cases.

By the decoupling property and (33), we have $V^\pi(S_1 - \{j\}) = V^\pi(S_1)$. Using $V^\pi(S_1 - \{j\}) = V^\pi(S_1)$ with (23), Lemma 2.2, and noting the optimality of $\tilde{\pi}$, we have

$$V^\pi(S_1 - \{j\}) = V^\pi(S_1) = V^{\tilde{\pi}}(S_1) \geq V^{\tilde{\pi}}(S_1 - \{j\}) \geq V^\pi(S_1 - \{j\}) \quad (34)$$

Relation (34) proves the induction step for (18) for $\pi(S_1) = i$.

This completes the induction step for (18). By induction, we have proved that (18) is true for π for all $S \subseteq \Omega$ where $|S| \geq 2$.

Finally, we prove that π is an optimal policy. First, note that by (18),

$$V^\pi(S) = V^{\tilde{\pi}}(S), \quad \forall S, |S| \geq 2 \quad (35)$$

Relation (18) also implies that for some $i \in \Omega$

$$V^\pi(\{j\}) = V^{\tilde{\pi}}(\{j\}), \quad \forall j \in \Omega \text{ s.t. } \pi(\{i\} \cup \{j\}) = j \quad (36)$$

We are left to consider $V^\pi(\{j\})$ when no such i as in (36) exists. By (21) we have $\pi(\{j\}) = \tilde{\pi}(\{j\})$ when $\pi(\{j\} \cup \{i\}) = i, \forall i \in \Omega$. Under this condition and (18), and identifying $S_1 = \{j\}$, π and $\tilde{\pi}$ satisfy the requirements of Lemma 2.3 (10) and (11). Lemma 2.3 (12) then implies that

$$V^\pi(\{j\}) = V^{\tilde{\pi}}(\{j\}), \quad \forall j \in \Omega \text{ s.t. } \pi(\{j\} \cup \{i\}) = i, \forall i \in \Omega \quad (37)$$

Equations (35), (36), and (37) prove the optimality of π . The proof of the lemma is complete. ■

Besides Lemma 2.4, we need one more result to prove Theorem 2.1, which is that the value function for Problem (P₁) is always an index function. We present this lemma next.

Lemma 2.5: For any optimal Markov policy $\tilde{\pi}$, $V^{\tilde{\pi}}(\cdot)$ is an index function on Ω .

Proof: First note that Lemma 2.2 implies Lemma 2.1 (3) is satisfied for $V^{\tilde{\pi}}(\cdot)$ on Ω .

Next, let π be the Markov policy satisfying (16) and (17) as constructed in Lemma 2.4. Consider any state S and $i \in S$ s.t. $V^\pi(S) \neq V^\pi(\{i\})$. If $\pi(S) = i$, then by removing all nodes except i from S via repeated application of Lemma 2.4 we would get $V^\pi(S) = V^\pi(\{i\})$, a contradiction. Hence, $\pi(S) \neq i$. So by Lemma 2.4 (18), we have $V^{\tilde{\pi}}(S) = V^{\tilde{\pi}}(S - \{i\})$. Thus the requirement (4) of Lemma 2.1 is satisfied for $V^{\tilde{\pi}}(\cdot)$ on Ω .

Since both requirements of Lemma 2.1 are satisfied, we have shown that $V^{\tilde{\pi}}(\cdot)$ is an index function on Ω . ■

We now use Lemma 2.1-Lemma 2.5 to prove Theorem 2.1.

Proof: [of Theorem 2.1]

Let π be the Markov policy satisfying (16) and (17) as constructed in Lemma 2.4. By Lemma 2.4 (3.), π is an optimal Markov policy. Hence, Lemma 2.5 indicates that $V^\pi(\cdot)$ is an index function.

We next show that π is a priority policy. By (16) and (17) we have

$$\pi(S) = i \implies \pi(S') = i, \quad \forall S' \subseteq S, i \in S' \quad (38)$$

$$\pi(S) = r_i \implies \pi(S') = r_i, \quad \forall S' \subseteq S, i \in S' \quad (39)$$

Properties (38) and (39) show that π is a priority policy (cf. Definition 2.4), with node priority as follows. For any S where $\pi(S) = i$ or $\pi(S) = r_i$, i has priority higher than all other nodes of S .

We have shown that π is a priority policy with index function $V^\pi(\cdot)$. Hence π satisfies Definition 2.6, and is an index policy. ■

Note that (38) and (39) imply that $\pi(S_1) = i$ and $\pi(S_2) = r_i$ cannot both occur for a given π for any S_1, S_2 with $i \in S_1$ and $j \in S_2$. That is, for a given system for Problem (P₁) and an optimal index policy π , if there is a state where π transmits from node i , then there is no state where π retires and receives the reward from i . Similarly, if there is a state where π retires and receives the reward from node i , then there is no state where π transmits at i . We henceforth use this fact when we write the node priority list for an index policy π , where exactly one of i and r_i is listed for each $i \in \Omega$. When r_i occurs in the list, this means that for states which include i and no nodes of higher priority than i , the optimal action is to retire. The reward received at retirement will be R_i .

It is interesting to note that the total number of stationary Markov policies for Problem (P₁) with N nodes is $N^{(2^N)}$, whereas the total number of priority policies is $N!$. Based on the result of Theorem 2.1, we develop an algorithm which is able to compute an optimal index policy for Problem (P₁) with computational complexity of only $O(N^2)$.

2) *Description of Centralized Algorithm:* As stated in Section II-A, we use the notation $V_i^\pi := V^\pi(\{i\})$.

Algorithm 1: (A Dijkstra-Type Algorithm for an Index Policy)

Define the sets \mathcal{A} and \mathcal{X} as follows.

Initially: \mathcal{A} contains the nodes of highest retirement reward R_{max} in arbitrary order (there must be at least one such node).

The action taken by the optimal index policy π on any node $i \in \mathcal{A}$ is r_i and the reward received is R_{max} . \mathcal{X} is the unordered complement (w.r.t. Ω) of \mathcal{A} .

During the construction of optimal policy π : \mathcal{A} contains a priority list of a set S of nodes, $S \subset \Omega$, together with the action specified by π on each node in S . \mathcal{X} is the unordered complement (w.r.t. Ω) of \mathcal{A} .

The algorithm proceeds as follows.

- 1) For each $i \in \mathcal{X}$, let π_i be an index policy with the same priority list as π for the nodes of \mathcal{A} , with i as the next highest priority node after \mathcal{A} , and with the priority of the nodes $\mathcal{X} - \{i\}$ arbitrary, but lower than i . Compute $V_i^{\pi_i}$ from

$$V_i^{\pi_i} = \max \left\{ \frac{-c_i + \sum_{S \supset \{i\}: \pi_i(S) \neq i} P^i(S|i) V_{\pi_i(S)}^{\pi_i}, R_i \right\} \quad (40)$$
- 2) Choose $i \in \mathcal{X}$ with the highest value of $V_i^{\pi_i}$, with ties broken arbitrarily. Append this node to the list \mathcal{A} as the next priority node, together with the action specified by (40). Remove i from \mathcal{X} .
- 3) If \mathcal{X} is empty, stop. If not, go to step 1.

Remark: In Step 1 the right-hand-side of (40) computes the best expected reward for node i , assuming i is the node of next highest index in π . This computation is feasible because π is a priority policy.

We now establish a relation between (40) and the optimality equation for Problem (P_1) . Such a relation allows us to prove that Algorithm 1 indeed computes an optimal index policy. Theorem 2.1 states that there is an optimal index policy for Problem (P_1) . Hence the Dynamic Programming equation for Problem (P_1) can be written $\forall i \in \Omega$

$$V_i^\pi = \max \left\{ \max_{\tilde{\pi}} \left\{ -c_i + \sum_{S \supset \{i\}} P^i(S|i) V_{\tilde{\pi}(S)}^{\tilde{\pi}} \right\}, R_i \right\} \quad (41)$$

where the inner maximum is taken over all index policies $\tilde{\pi}$. In the following lemma we use the existence of optimal index policy π to put the computation of V_i^π into a more convenient form.

Lemma 2.6: Assume π is an optimal index policy for Problem (P_1) . Then $\forall i \in \Omega$

$$V_i^\pi = \max \left\{ \max_{\tilde{\pi}} \left\{ \frac{-c_i + \sum_{\mathcal{N}(i) \supset S \supset \{i\}: \tilde{\pi}(S) \neq i} P^i(S|i) V_{\tilde{\pi}(S)}^{\tilde{\pi}}}{\sum_{\mathcal{N}(i) \supset S \supset \{i\}: \tilde{\pi}(S) \neq i} P^i(S|i)}, R_i \right\} \right\} \quad (42)$$

where the inner maximum is taken over all index policies $\tilde{\pi}$.

Proof: The proof of Lemma 2.6 is in Appendix A. ■

In the following corollary we show an important property of the update (40) and its relation to (41).

Corollary 2.1: Assume π is an optimal index policy for Problem (P_1) . Let π_i be as defined in Algorithm 1, that is, π_i is the same as π for the set of highest priority nodes \mathcal{A} of π , and node i is the node of highest priority in \mathcal{X} according to π_i . Assume i is also the node of highest priority in \mathcal{X} according to π . If $V_i^{\pi_i}$ is computed as in (40), then

$$V_i^{\pi_i} = V_i^\pi \quad (43)$$

Proof: Corollary 2.1 follows from the fact that (40) computes (42) with the policy $\tilde{\pi} = \pi_i$, which is optimal for all nodes of priority i or higher. ■

Algorithm 1 also resembles Klimov's algorithm [14] and has the following feature.

Theorem 2.2: For Problem (P_1) , Algorithm 1 produces an optimal index policy.

Proof: We prove the theorem by induction on the number of nodes in the set \mathcal{A} defined in the description of Algorithm 1. Recall that for a Markov decision problem, the optimal policy maximizes the value function for each state.

Let π be an optimal index policy. Suppose Algorithm 1 has run to the point that $|\mathcal{A}| = L$. Let $i \in \mathcal{X}$ be the node with the actual $(L+1)$ 'th highest priority according to π , whether retiring or not. Let $j \in \mathcal{X}, j \neq i$. Let π_j denote the priority policy that has the same priority as π in the first L nodes, gives j the $(L+1)$ 'th node priority, retiring or not as optimal, and arbitrarily gives priority lower than $L+1$ to the remaining nodes.

Then we claim that

$$V_i^{\pi_i} = V_i^\pi \geq V_j^\pi \geq V_j^{\pi_j} \quad (44)$$

The equality of (44) follows from Corollary 2.1, because i is assumed to be the actual $(L+1)$ 'th priority node of π . The first inequality of (44) follows because by assumption i is higher priority than j in π . The second inequality follows because π is an optimal policy.

Relation (44) implies that any node $j \in \mathcal{X}$ maximizing $V_j^{\pi_j}$ may optimally be made the $(L+1)$ 'th priority node. Note that this j is not necessarily unique. This is the procedure used to find the node of next highest priority in Step 2 of Algorithm 1.

This completes the proof of the induction step. Hence, by induction Algorithm 1 produces an optimal index policy. ■

3) *Remarks:*

- 1) In typical network models the network layer requests point-to-point communication from the link layer. However, the stochastic local broadcast nature of our link model requires a control which describes alternative courses of action depending on feedback from random transmission events. In light of Theorem 2.1, we can state precisely what this means for Problem (P_1) . The instructions for transmission at a node consist of a priority list of neighbors and the transmitting node itself. The node transmits until a node of higher priority successfully receives the packet. Note that there is no *one* route, as the actual route a packet takes between source and destination is sample path dependent.
- 2) The index policy result of Theorem 2.1 is very general. It says that no matter what the actual values of $(P, c, R)_i$ are at each node, an index policy is always one possible optimal policy *structure*. There might be errors in knowledge of $(P, c, R)_i$, and this affects the indices of the computed optimal index policy. But the fact that *some* index policy is optimal is always true.
- 3) We have shown that Algorithm 1 computes an optimal index policy for Problem (P_1) . It is important to note that Algorithm 1 uses all of the network parameters, meaning the $(P, c, R)_i$ values at each node. To run

this algorithm, all of this information must be available at the same location. In this sense, Algorithm 1 is a centralized algorithm, like the Dijkstra algorithm used in OSPF (see [15] Ch. 4). The number of updates of (40) in Algorithm 1 is in the worst case $O(N^2)$.

4) Distributed Implementation of an Optimal Index Policy:

We note an interesting feature of an index policy used for Problem (\mathbf{P}_1) . The indices of the network nodes are fixed, and at each transmission a node of highest index which has the packet is chosen to transmit. This leads to the following property.

Property 2.1: In an index policy for Problem (\mathbf{P}_1) , the only nodes of index higher than the transmitting node which can receive the packet are neighbors of the transmitting node.

Property 2.1 allows for a natural distributed implementation of the index policy, as follows. Imagine there is a token associated with the packet that begins with the packet at the node of origin. The token indicates which node is to transmit next. After a node transmits, it passes the token to a neighbor (that depends on the outcome of the transmission), or keeps the token for retransmission. By Property 2.1, an optimal index policy can be implemented in this way, as the optimal next node is always a neighbor. Note that there is no central control of this token passing mechanism. All decisions are made locally, and involve only neighboring nodes.

It is important to distinguish distributed optimal routing policy implementation from distributed computation of an optimal index policy itself. For policy implementation, it is assumed that the index policy has already been determined. We consider the problem of distributed index policy computation starting in Section IV.

III. THE STOCHASTIC ROUTING PROBLEM WITH TRANSMISSION CONTROL

In this section we extend the model of Problem (\mathbf{P}_1) to allow for control of transmission type at each node. That is, at each time step the controller may choose a node for transmission, and also a type of transmission at that node. Transmission type may be used to model various physical layer features, such as multiple transmission power levels, modulation/coding scheme, antenna directionality, and destination addressing. We begin with some notation and definitions.

A. Notation and Definitions

Definition 3.1: W_i refers to the number of transmission types available at node i .

Definition 3.2: We write $\pi(S) = (i, k)$ to mean that when in state S , policy π chooses node i and transmission type k , $i \in \Omega, k \in \{1 \dots W_i\}$. The expression $\pi(S) = (i, *)$ means policy π chooses i at some unspecified transmission type. The notation for retirement $\pi(S) = r_i$ is retained unchanged.

Definition 3.3: When in state S for which $\pi(S) = (i, k)$, a cost $c_{(i,k)}$ is incurred, and the transition probabilities are written $P^{(i,k)}(S'|S)$.

B. Statement of Problem

Problem (\mathbf{P}_2)

We consider Problem (\mathbf{P}_1) , with the following addition. At each time step the central controller chooses a node for transmission, from among the nodes with the packet, and a transmission type from among the allowable types for that node. To each node and transmission type is associated a transmission cost and a probability that a given set of nodes receives the packet. We seek a policy which maximizes (2) under the conditions of Problem (\mathbf{P}_1) and the above addition.

C. Analysis of Problem (\mathbf{P}_2)

The time-homogenous Markov nature of Problem (\mathbf{P}_1) is not altered by adding in a choice of transmission type. We find that with an appropriate mapping to a new larger space (each transmission type becomes a node), we can apply the result of Problem (\mathbf{P}_1) directly. To demonstrate this, we define a new problem, Problem (\mathbf{P}_3) , show its relation to Problem (\mathbf{P}_2) , and then show that it is equivalent to Problem (\mathbf{P}_1) .

1) *Notation and Definitions for Problem (\mathbf{P}_3) :* We define a new space of nodes Ω_P as follows. As in Definition 3.2, we list possible control choices for Problem (\mathbf{P}_2) as (i, k) , where $i \in \Omega$ and $k \in \{1 \dots W_i\}$. Every node of Ω_P corresponds to exactly one such control choice of Problem (\mathbf{P}_2) . Defining N_P to be the cardinality of Ω_P , we then have

$$N_P := |\Omega_P| = \sum_{i=1}^N W_i \quad (45)$$

where as before $N := |\Omega|$. We write $j \in \Omega_P$ to refer to a node from Ω_P . We can also refer to the (i, k) pair of a node $j \in \Omega_P$, which is the associated i and transmission type $k \in \{1 \dots W_i\}$ in the original space Ω . We say that the group of nodes in Ω_P which corresponds to $i \in \Omega$ is the *family* of i , so that the family of i has W_i members. The family of $S \subseteq \Omega$ is the set of all nodes in Ω_P which are in the family of any one of the nodes of S .

The cost to transmit from a node j with (i, k) pair is defined to be $c_{(i,k)}$, the cost for transmission type k at i in Problem (\mathbf{P}_2) .

Transmissions in the Ω_P space are based on the corresponding events in the Ω space, as follows. Say a node j with (i, k) pair is chosen for transmission. This incurs a cost $c_{(i,k)}$. On Ω , transmission from i with transmission type k leads to a set of nodes, say $S_1 \subseteq \Omega$, receiving the packet. Correspondingly, on Ω_P by definition a node $j \in \Omega_P$ with (l, k) pair receives the packet if and only if l receives the packet on Ω , $l \in S_1$. That is, nodes in Ω_P in the family of l receive the packet precisely when l does. Note that this means that packet reception for nodes of Ω_P in the same family are deterministically coupled, in that they all receive or all do not receive it. This strongly restricts the kinds of transitions that can occur in Ω_P .

Each member of the family of $i \in \Omega$ gets the same reward as i . That is, let R_P be the reward function on Ω_P , an index function on Ω_P . If $j \in \Omega_P$ is in the family of i , then $R_P(\{j\}) = R(\{i\})$. Because R_P is an index function, this fully defines R_P on Ω_P .

2) Formulation of Problem (\mathbf{P}_3):

Problem (\mathbf{P}_3) is a formulation of Problem (\mathbf{P}_1) for the space Ω_P , with the transition probabilities, cost, and reward described above. The starting state is the family of S_o , the initial state for Problem (\mathbf{P}_2).

At each time point the controller either chooses a node from Ω_P for transmission, or retires. We seek a policy which maximizes (2) under the conditions of Problem (\mathbf{P}_1) defined on Ω_P .

3) *Relation of Problem (\mathbf{P}_3) to Problem (\mathbf{P}_2):* We show the relation of Problem (\mathbf{P}_3) to Problem (\mathbf{P}_2). There is a one-to-one mapping from the states of Ω -space to the states of Ω_P -space, as follows. Let $S_1 \subseteq \Omega$ be a state of Problem (\mathbf{P}_2). The state of Problem (\mathbf{P}_3) corresponding to S_1 is the set of all nodes $j \in \Omega_P$ such that j is in the family of some node $i \in S_1$. There are states of Problem (\mathbf{P}_3) which do not correspond to any state in Problem (\mathbf{P}_2). To see this, let $j, k \in \Omega_P$ both be in the family of $i \in \Omega$. Consider a state $S \subset \Omega_P$ where $j \in S, k \notin S$; this state S is not the mapping from any state $S_1 \subseteq \Omega$. The one-to-one mapping from Ω to Ω_P is not in general an *onto* mapping, because $|\Omega_P| > |\Omega|$. Those states $S \subseteq \Omega_P$ that are not the image of any state $S_1 \subseteq \Omega$ under the above-mentioned mapping are not reachable under the state transition mechanism of Problem (\mathbf{P}_3), and play no role in the analysis.

A control action for Problem (\mathbf{P}_3) has a corresponding action in Problem (\mathbf{P}_2), as follows. A node $j \in \Omega_P$, with associated (i, k) , chosen for transmission in Problem (\mathbf{P}_3) corresponds to transmission at node i of type k in Problem (\mathbf{P}_2). The cost $c_{(i,k)}$ incurred for this transmission is the same in both problems, and the state reached by the transmission in Ω maps to the state reached in Ω_P . When retirement is chosen for Problem (\mathbf{P}_3), this corresponds to retirement for Problem (\mathbf{P}_2), and the same reward is received, because corresponding states have the same reward.

Hence, Problem (\mathbf{P}_2) and Problem (\mathbf{P}_3) are entirely equivalent, in that to each decision policy for Problem (\mathbf{P}_3) there is a corresponding decision policy for Problem (\mathbf{P}_2) which results in exactly the same behavior, and hence expected reward, for both systems. We can thus solve Problem (\mathbf{P}_2) by finding an optimal policy for Problem (\mathbf{P}_3). We proceed to solve Problem (\mathbf{P}_3).

4) *Analysis of Problem (\mathbf{P}_3):* We show that the system of Problem (\mathbf{P}_3) satisfies the requirements of Problem (\mathbf{P}_1), and hence is a special case of that system. Specifically, we show that the increasing and decoupling properties (cf. Definition 2.1 and Definition 2.2) are satisfied, and that the reward function is an index function.

The increasing property on Ω_P follows directly from the fact that the increasing property holds on the underlying space Ω , together with the way states of Ω map to Ω_P . That is, transmissions only lead to an increasing state in Ω , leading to an increasing state in Ω_P .

The decoupling property holds for Problem (\mathbf{P}_3) because it holds for Problem (\mathbf{P}_2). The nodes that receive transmissions in Problem (\mathbf{P}_2) are unaffected by what other nodes have the packet.

Note that transmission events in Ω_P can be highly corre-

lated, in that nodes of Ω_P in the same family either all have the packet, or none have it. However, such event correlations are allowed by Model (\mathbf{M}).

Finally, R_P satisfies the definition of an index function on Ω_P , because R is an index function on Ω and R_P gives the same reward for the associated states of Ω_P .

Hence, an index policy is optimal for Problem (\mathbf{P}_3), and Algorithm 1 can be used to determine such a policy. This policy is also optimal for Problem (\mathbf{P}_2). Thus, Algorithm 1 is effectively used to solve Problem (\mathbf{P}_2).

Note that in the resulting priority list of an optimal index policy π , nodes from the same family in Ω_P will appear in some relative order. Since nodes in a given family either all have the packet or none have it, it is clear that only one node from each family will ever be used for transmission by the algorithm. Hence, all the nodes that are not of the highest priority within their family can be removed from the priority list for simplicity. In the terms of Problem (\mathbf{P}_2), this means that only one transmission type at each node is ever used for transmission.

Algorithm 1 when applied to Problem (\mathbf{P}_2) in this manner is $O(N_P^2)$ in complexity, that is, of complexity $O\left(\left(\sum_{i=1}^N W_i\right)^2\right)$.

Comment: The results in this section can be used to show that when a choice exists among deterministic transmissions to single neighbor nodes, then Algorithm 1 reverts to the well-known Dijkstra's Algorithm ([15] Ch. 4). In this sense Algorithm 1 is a generalization of Dijkstra's algorithm to the case of stochastic local broadcast transmissions.

IV. NOTATION USED IN DISTRIBUTED ALGORITHMS FOR PROBLEM (\mathbf{P}_1)

A. New Notation and Definitions

We summarize the notation and definitions we need for the distributed algorithms in this paper.

A general Markov policy can be written

$$\{\pi_1 \pi_2 \dots \pi_t \dots\} \quad (46)$$

where the subscript of π indexes time. Since Problem (\mathbf{P}_1) is a time-homogeneous Markov decision problem, we know there exists an optimal stationary policy of the form

$$\{\pi \pi \dots \pi \dots\} \quad (47)$$

Further, from Theorem 2.1 we know that there exists a stationary optimal policy of the form (47) where π is an index policy.

When describing distributed algorithms to compute an optimal policy for Problem (\mathbf{P}_1), we need to consider a more general type of policy.

Definition 4.1: A *local index policy* for node i at t is written π_t^i , and defines a node ordering of $\mathcal{N}(i)$, $i \in \Omega$, with retirement for node i indicated by r_i if desired. A *distributed index policy* at t is written $\Pi_t = \{\pi_t^1 \pi_t^2 \dots \pi_t^N\}$, where π_t^i is a local index policy for node i . A *distributed index policy* is written $\Pi = \{\Pi_1 \Pi_2 \dots \Pi_t \dots\}$.

A distributed index policy functions by transmitting at the current node, say i , at t , then using the node ordering π_t^i to choose the next node for transmission.

For local index policy π_t^i , we write $j >_{\pi_t^i} k$ when j has higher priority than k under π_t^i , $j, k \in \mathcal{N}(i)$.

Definition 4.2: Consider nodes i and j , local index policies π_t^i and π_t^j , and let $k, l \in \mathcal{N}(i) \cap \mathcal{N}(j)$. If either 1) $k >_{\pi_t^i} l$ and $k >_{\pi_t^j} l$ or 2) $k <_{\pi_t^i} l$ and $k <_{\pi_t^j} l$, then we say π_t^i and π_t^j *match* on k and l .

Definition 4.3: If π_t^i and π_t^j match on k and $l \forall i, j \in \Omega, \forall k, l \in \mathcal{N}(i) \cap \mathcal{N}(j)$, we say distributed index policy Π_t is *uniform* at t .

When Π_t is uniform at t , a global order of the nodes is induced. We call the index policy which has this global node ordering the *associated index policy* of Π_t .

We write $P^i(S|i)$ to indicate the transmission probability as known at node i .

V. A DISTRIBUTED ALGORITHM FOR PROBLEM (P₁)

We present our first distributed algorithm (Algorithm 2 described below) which computes the optimal solution for Problem (P₁), and has the characteristic that computations at each node use only information directly from neighbor nodes. This property is critical to the distributed implementation of an optimal policy in an ad hoc wireless network. We claim convergence of the algorithm to the global node ordering and value function consistent with the optimal index policy under the following constraints.

- 1) Each node i keeps a current estimate, denoted by V_i , of its own optimal expected reward value, with initial value $0 \leq V_i^0 \leq R_{max}$. Each node i also stores the most recently received estimate of each of its neighbors' optimal expected reward values, denoted $V_{i,j}$, where $j \in \mathcal{N}(i)$, with initial value $0 \leq V_{i,j}^0 \leq R_{max}$.
- 2) Information transfer among neighboring nodes consists only of the current V_i value of the transmitting node.
- 3) Each node's V_i information is transmitted asynchronously.
- 4) A node's V_i update, defined below, is also asynchronous.
- 5) It is assumed that each node has knowledge of its own $P^i(S|i)$ update structure. For example, node i may estimate $P^i(S|i)$ based on all its communications, both control signals and packets.
- 6) The energy required to run the algorithm is not included in finding the optimal solution for Problem (P₁).

The algorithm is as follows.

Algorithm 2: An event time n is when one or more of the following two events occurs. Any number of events may occur at an event time.

Event 1 A node i receives V_j from a neighbor j and stores it as $V_{i,j}^n$, $j \in \mathcal{N}(i)$.

Event 2 A node i recomputes V_i using the current $V_{i,j}$ values, as follows.

$$V_i^n = \max \left\{ \max_{\tilde{\pi}} \left\{ \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i) V_{i, \tilde{\pi}(S)}^n}{\sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i)} \right\}, R_i \right\} \quad (48)$$

The maximization in (48) is over all local priority policies $\tilde{\pi}$ of i .

It is assumed that events 1 and 2 occur infinitely often at each node i .

Note that $P^i(S|i) > 0$ and $\tilde{\pi}(S) \neq i$ imply that $S \supset i$, and for conciseness in (48) we write simply $S \subseteq \mathcal{N}(i)$.

We require no *a priori* time ordering on the above events, nor on the nodes where they are occurring. At times when neither of the above events is taking place, the system is in a frozen state, with all system parameters remaining unchanged. An event which occurs at some event time can have no effect on other events at the same time. Hence, we can choose an arbitrary order for all events occurring at a given time without affecting the outcome. In this way, we can talk sensibly about the n 'th event in the system since the start, and we use this convention hereafter.

The local policy which optimizes (48) for node i at event n is a local index policy at i , $\pi_{t_n}^i$. For convenience we notate this as π_n^i , where the context prevents ambiguity. The distributed index policy after event n will be denoted by Π_n .

Let V_i^n be the expected reward for node i just after event n in the above system, so that V_i^0 is this value at the start of the algorithm, where the allowed range is $0 \leq V_i^0 \leq R_{max}$. Let $V_{i,j}^n$ be the value after event n of the last transmitted V_j^m received at i , where $m < n$ is the event index of this transmission. At the start, the $V_{i,j}^0$ values do not need to match the neighbor's V_j^0 values, we only require that $0 \leq V_{i,j}^0 \leq R_{max}$.

The computation of (48) is based on update (40) of Algorithm 1. Given the nature of this update equation, finding the maximum in (48) is easier than it might first appear.

A Method for Finding the Maximum Over $\tilde{\pi}$ in (48)

Method 1: Let event n be i 's computation of V_i^n . Our goal is to compute

$$H := \max_{\tilde{\pi}} \left\{ \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i) V_{i, \tilde{\pi}(S)}^n}{\sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i)} \right\} \quad (49)$$

When i has no neighbors, we define $H = -\infty$, and we have $V_i^n = R_i$. If node i has at least one neighbor, proceed as follows.

- 1) Rank order the values $V_{i,j}^n$ that i has most recently received from its neighbors.

Let $\xi_1, \xi_2, \dots, \xi_l$ indicate the subscripts of i 's neighbors in rank order, where l is the number of i 's neighbors. That is, ξ_j is the subscript of the j 'th best neighbor of i according to the ranking of $V_{i,j}^n$. Ties are decided arbitrarily. Denote by k the rank of node i w.r.t. its neighbors. Denote by $\tilde{\pi}_k$ the priority policy which ranks the k highest neighbors of node i above i . Initially set $k = 2$. (k is a dummy variable in this computation).

- 2) Compute

$$W = \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}_k(S) \neq i} P^i(S|i) V_{i, \tilde{\pi}_k(S)}^n}{\sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}_k(S) \neq i} P^i(S|i)} \quad (50)$$

- 3) If $V_{i, \xi_{k-1}}^n > W \geq V_{i, \xi_k}^n$, then set $H = W$ and halt (we have found the rank of node i w.r.t. its neighbors).
- 4) If $k \leq l$, modify $\tilde{\pi}_k$ by moving neighbor ξ_k to higher priority than i , leaving all else in $\tilde{\pi}_k$ unchanged. Set $k = k + 1$, and go to 2.
- 5) If $k > l$, then node i is the worst node of all the neighbors. Set $H = W$ and halt.

The process continues until either the condition of step 3, or step 5, is satisfied. Termination is guaranteed, since the condition of step 5, must eventually be satisfied.

Note that the denominator in (50) is not zero as long as i has at least one neighbor.

Lemma 5.1: Method 1 finds a local policy $\hat{\pi}$ which maximizes (49).

Proof: Method 1 can be viewed as an implementation of Algorithm 1 of Section II for $\Omega = \mathcal{N}(i)$ where $R_i = -\infty$, and where for all nodes $j, k \in \mathcal{N}(i), j \neq i$, we have

- 1) $P_{jk} = 0$
- 2) $R_j = V_{i,j}^n$

Under these conditions, each V computation (given by (48)) for node $j \neq i$ gives $V_{i,j}^n$, and Algorithm 1 of Section II reduces to Method 1. Since Algorithm 1 of Section II finds the optimal policy for all nodes, this policy maximizes (49). ■

Comments

- 1) Using Method 1 to compute (49) is in the worst case $O(l^2)$, whereas an exhaustive search over all local policies $\hat{\pi}$ (i.e. local node rankings) would be $O(l!)$.
- 2) In a practical implementation, there may be ways to speed up Method 1 by changing the order in which the neighbor nodes are put into or removed from the list of better nodes. For example, we might want to always try the order which worked at the last update, and then add or remove nodes sequentially from there.

We proceed with the analysis of Algorithm 2. The following theorem is our main result concerning Algorithm 2, summarizing its convergence properties.

Theorem 5.1: For Algorithm 2 with any initial state s.t. $0 \leq V_i^0 \leq R_{max}$ and $0 \leq V_{i,j}^0 \leq R_{max}$ for all $i, j \in \Omega$, we have

$$\lim_{n \rightarrow \infty} V_i^n = V_i^\pi, \forall i \in \Omega \quad (51)$$

- 2) There exists an event $n_p < \infty$ s.t. $\Pi = \{\Pi_{n_p} \Pi_{n_p+1} \Pi_{n_p+2} \dots\}$ is an optimal distributed policy
- 3) If $V_i^\pi \neq V_j^\pi, \forall i \in \Omega, j \in \mathcal{N}(i)$ (52)

a) $V_i^n = V_i^\pi, \forall n \geq n_p, \forall i \in \Omega$

b) $\Pi = \{\Pi_{n_p} \Pi_{n_p+1} \Pi_{n_p+2} \dots\}$ has an associated index policy that is optimal

Before proving Theorem 5.1, we present a series of lemmas which we make use of in the proof. Our method of proof of Theorem 5.1 is influenced by the proof of the Asynchronous Distributed Bellman-Ford algorithm in [4]. We begin with a simple proof of an important property of the node update procedure which is fundamental for the analysis.

Lemma 5.2 (Update Monotonicity): Consider two cases of a node recomputation for i at event n using (48). In case 1, the neighbor values are $\tilde{V}_{i,j}^n$ and the updated node value is \hat{V}_i^n . In case 2 the neighbor values are $\hat{V}_{i,j}^n$ and the updated node value is \hat{V}_i^n . Assume $\hat{V}_{i,j}^n \geq \tilde{V}_{i,j}^n, \forall j \in \mathcal{N}(i)$. Then

$$\hat{V}_i^n \geq \tilde{V}_i^n \quad (53)$$

Proof: For a given $\hat{\pi}$, (50) is monotonic in each $V_{i,j}^n$. Let $\hat{\pi}$ be the policy chosen in (49) for case 1, and let $\hat{\pi}$ be the policy chosen in (49) for case 2. Let \hat{H} be the value in (50)

when case 2's policy $\hat{\pi}$ is used with case 1's values $\tilde{V}_{i,j}^n$. Then

$$\hat{H} \leq \bar{H} \leq \hat{H} \quad (54)$$

where the final inequality follows because $\hat{\pi}$ is optimal in (49).

Equation (48) and inequality (54) imply (53). ■

We say that the *monotonicity property* holds for a node update due to the result of Lemma 5.2.

We now define two random sequences, \underline{D}_i^n and \bar{D}_i^n , which we use to provide bounds on node updates. Noting that at the start of Algorithm 2 each node has a V value between 0 and R_{max} , and aiming at a kind of worst case initial state in light of the monotonicity property just demonstrated, we define \underline{D}_i^n and \bar{D}_i^n as follows.

Definition 5.1: We define $\forall i \in \Omega$ such that $R_i \neq R_{max}$

$\underline{D}_i^n :=$ Computed value for i after event n when $V_i^0 = V_{i,j}^0 = 0$

$\bar{D}_i^n :=$ Computed value for i after event n when $V_i^0 = V_{i,j}^0 = R_{max}$

For i such that $R_i = R_{max}$, define

$$\underline{D}_i^n = \bar{D}_i^n := R_{max}, \forall n$$

This last definition results from the fact that a node which can retire and receive R_{max} should always do so.

In the following, we assume π is an optimal policy, and that V_i^π is the optimal value function for i .

In the next lemma, we demonstrate that \underline{D}_i^n is a monotonically non-decreasing sequence which lower bounds V_i^n , and in turn is upper bounded by V_i^π . Similarly, we demonstrate that \bar{D}_i^n is a monotonically non-increasing sequence which upper bounds V_i^n , and in turn is lower bounded by V_i^π . Once these facts are proved, it will remain to show that \underline{D}_i^n converges to V_i^π from below, and that \bar{D}_i^n converges to V_i^π from above, to obtain convergence of V_i^n to V_i^π .

Lemma 5.3: We have

- 1) $\underline{D}_i^m \leq \underline{D}_i^n \leq V_i^n \leq \bar{D}_i^n \leq \bar{D}_i^m, \forall i, n, m \leq n$
- 2) $\underline{D}_i^n \leq V_i^\pi \leq \bar{D}_i^n, \forall i, n$

Proof: We proceed by induction on event number. Consider the following set of equations at some event n .

- i) $\underline{D}_i^l \leq \underline{D}_i^m \leq V_i^m \leq \bar{D}_i^m \leq \bar{D}_i^l, \forall i, l \leq m \leq n$
- ii) $\underline{D}_i^n \leq V_i^\pi \leq \bar{D}_i^n, \forall i$

At $n = 0$, we have the starting values $\underline{D}_i^0 = 0$ and $\bar{D}_i^0 = R_{max}$. Hence (ii) is true. Also, the starting value for the actual system must be between these extreme values $0 \leq V_i^0 \leq R_{max}$ for any node i , so (i) is true at $n = 0$ also.

Assume that after event n (i) and (ii) are true. We want to prove that (i) and (ii) are true after event $n + 1$. We consider the two possible event types at $n + 1$, as follows.

- 1) A node transmits its V value. This has no effect on (i) and (ii) above, and they remain valid for $n + 1$.
- 2) A node (call it i) recomputes its V value. Since all nodes other than i are unaffected, we can focus on i at $n + 1$. The last time i 's V value was computed, some estimate of each neighbor's V value was used. Consider neighbor j , and let l be the past event at which j computed the V value which was used by i for j at its last V value computation. Similarly, let m be the past event at which j computed its value which is being used by i for the

current computation at $n + 1$. Then clearly $l \leq m \leq n$, and by (i) and the induction hypothesis we have

$$\underline{D}_j^l \leq \underline{D}_j^m \leq V_j^m \leq \overline{D}_j^m \leq \overline{D}_j^l \quad (55)$$

Since (55) is true for all neighbors of i , by the monotonicity property (Lemma 5.2) we have

$$\underline{D}_i^n \leq \underline{D}_i^{n+1} \leq V_i^{n+1} \leq \overline{D}_i^{n+1} \leq \overline{D}_i^n \quad (56)$$

Equation (56) and our assumption of (i) at n together imply that (i) is also valid for i at $n + 1$.

Next, by inductive assumption (ii) we have for any neighbor j at event m , where m is as above,

$$\underline{D}_j^m \leq V_j^\pi \leq \overline{D}_j^m \quad (57)$$

By Lemma 2.6, if each neighbor j has its correct optimal value V_j^π , then (48) for i gives V_i^π , the optimal value for i . So again monotonicity (Lemma 5.2) implies that

$$\underline{D}_i^{n+1} \leq V_i^\pi \leq \overline{D}_i^{n+1} \quad (58)$$

and (ii) is also true at $n + 1$.

Hence for both types of events (i) and (ii) remain true at $n + 1$, and the induction step is complete. ■

Comment: In Lemma 5.3, the first part follows directly from the monotonicity property, and in fact is true for any function that is computed at each node that has this property. The second part further requires the property that when all neighbors are correct, the correct value function gets computed.

Corollary 5.1: We have

- 1) $\underline{D}_i^n = V_i^\pi \implies \underline{D}_i^m = V_i^\pi, \forall m \geq n$
- 2) $\overline{D}_i^n = V_i^\pi \implies \overline{D}_i^m = V_i^\pi, \forall m \geq n$

Proof: These two results follow directly from Lemma 5.3. ■

In the following lemma, we demonstrate convergence of \underline{D}_i^n to the optimal value V_i^π in a finite number of steps (we remind the reader that π denotes an optimal policy).

Lemma 5.4: There exists $n_p < \infty$ s.t. $\forall n \geq n_p, \underline{D}_i^n = V_i^\pi, \forall i \in \Omega$.

Proof: The proof of Lemma 5.4 is in Appendix B. ■

Unfortunately, finite time convergence does not hold in general for \overline{D}_i^n , as the following example demonstrates.

Example

Consider the system of Figure 1 with parameters $0 < p < 1$, $0 < q \leq 1$, $c_1 = c_2 = 1$, and $R_1 = R_2 = 0, R_3 = R$. That is, Q_3 is the destination node. Assume that $R > \frac{1}{p}$, so the system has a non-trivial optimal policy. Transmission success from either Q_1 or Q_2 to the other two nodes is independent, with probabilities p and q respectively. Assume Algorithm 2 begins with $V_i^0 = V_{i,j}^0 = R, \forall i, j$.

As Algorithm 2 runs, the node value computations in Q_1 and Q_2 ping-pong back and forth, as the update of one is transmitted to be used in the other's update. Letting V_1^n denote the update value at Q_1 at the n th such update (i.e. the n th update where V_1 actually changes value), it can be shown that

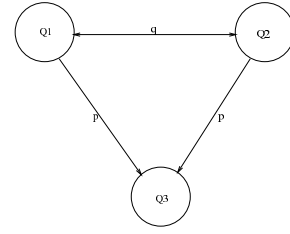


Fig. 1. System which takes infinite time to converge

for $n \geq 1$,

$$\begin{aligned} V_1^n &= \frac{(1-p)^2 q^2}{(p + (1-p)q)^2} V_1^{n-1} + \frac{(pR - 1)(p + 2(1-p)q)}{(p + (1-p)q)^2} \\ &:= A V_1^{n-1} + B \end{aligned} \quad (59)$$

where $V_1^0 = R$, and A and B are defined by (59). Since $0 < A < 1$, standard difference equation methods [35] yield the closed form solution

$$V_1^n = \left(R - \frac{B}{1-A} \right) A^n + \frac{B}{1-A} \quad (60)$$

A similar equation holds for V_2^n . Since $A < 1$ and $\frac{B}{1-A} = R - \frac{1}{p}$, we obtain $\lim_{n \rightarrow \infty} V_1^n = R - \frac{1}{p}$, which is the correct value function. But because $A > 0$, this value is never reached in finite time. □

Below we discuss the sense in which an optimal distributed policy is reached in finite time in the above example.

Comment: The key fact in the preceding example is that nodes 1 and 2 have identical value functions. As a result, the V_i^n values of these two nodes converge to the same value, and hence interact in each computation at each event n , preventing convergence in finite time. When limit values are different at each node, a result similar to Lemma 5.4 for the \overline{D}_i^n sequence follows easily, as shown in the following lemma.

Lemma 5.5: For Algorithm 2,

- 1) $\forall i$, there exists $W_i \in \mathbb{R}$ s.t. $\overline{D}_i^n \downarrow W_i$.
- 2) Assume $W_i \neq W_j, \forall i, j$ with $j \neq i$. Then $\exists n_p < \infty$ s.t. $\forall n \geq n_p, \overline{D}_i^n = V_i^\pi, \forall i$.

Proof: The proof of Lemma 5.5 is in Appendix C. ■

We are able to prove asymptotic convergence for the \overline{D}_i^n sequence in general. To accomplish this, we define a new system, which we refer to as the *round-robin* (r-r) system. In the r-r system, node update and transmission events follow a fixed pre-defined order. The r-r system is more analytically tractable than the general asynchronous system of Algorithm 2. Using the monotonicity property (Lemma 5.2), we are able to prove that \overline{D}_i^n is bounded above by a corresponding value in the r-r system (specified in Definition 5.3), and then show that they both must converge to V_i^π . We begin with the basic definitions.

Definition 5.2: The *round-robin* (r-r) system is defined as the system of Problem (\mathbf{P}_1) restricted to the following order of events. For each $1 \leq i \leq N$, node i computes V_i , successfully transmits V_i to its neighbors, and then these events are repeated at $(i+1) \bmod N$. For consistency in the following, we assume the system starts at $i = 1$.

Definition 5.3: For the r-r system:

- Define u_n^i to be the event number of the n th computation of the V value of i .
- Let W_i^k be the V value at u_n^i for i with initial state R_{max} for all nodes in the system.

For the actual system:

- Define v_n^i to be the event number of the first computation of the V value at i after each node j , $j < i$ has completed the n th computation of its V value, and each node j , $j > i$ has completed the $(n-1)$ th computation of its V value.

Note that the u_n^i 's and v_n^i 's are all finite, since we have assumed that each of these events occurs infinitely often. From the above definitions, we immediately infer that

$$v_n^i \geq u_n^i, \quad \forall i, n \quad (61)$$

We show that $W_i^{u_n^i}$ bounds $\bar{D}_i^{v_n^i}$.

Lemma 5.6: We have

$$\bar{D}_i^{v_n^i} \leq W_i^{u_n^i}, \quad \forall i, n \quad (62)$$

That is, the n th computation of the V value at i in the actual system is bounded above by that for the r-r system, when both start all nodes at R_{max} , for all i and number of computations n .

Proof: The proof of Lemma 5.6 is in Appendix D. ■

Since Lemma 5.3 is valid for any event order, the results there are also valid for the r-r system. Hence for each i , W_i^k is a non-increasing sequence which is bounded below by V_i^π . Hence it converges to a limit (see [30] Thm 3.14). We will show that this limit must be V_i^π , so that in general $W_i^k \downarrow V_i^\pi, \forall i$.

Definition 5.4: We write the update equation for i in the r-r system, and define the update function f_i as follows.

$$\begin{aligned} W_i^{k+1} &= \max \left\{ \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \pi_{k+1}(S) < i} P^i(S|i) W_{\pi_{k+1}(S)}^{k+1}}{\sum_{S \subseteq \mathcal{N}(i): \pi_{k+1}(S) \neq i} P^i(S|i)} \right. \\ &\quad \left. + \frac{\sum_{S \subseteq \mathcal{N}(i): \pi_{k+1}(S) > i} P^i(S|i) W_{\pi_{k+1}(S)}^k}{\sum_{S \subseteq \mathcal{N}(i): \pi_{k+1}(S) \neq i} P^i(S|i)}, R_i \right\} \\ &:= \max \{f_i(W_1^{k+1}, W_2^{k+1}, \dots, W_{i-1}^{k+1}, W_{i+1}^k, \dots, W_N^k), R_i\} \end{aligned} \quad (63)$$

For convenience, we will not notate the k 's in (63) when writing the formula for f_i , since they are fixed. Define

$$\mathbf{W}^{k,i-} := (W_1^{k+1}, W_2^{k+1}, \dots, W_{i-1}^{k+1}, W_{i+1}^k, \dots, W_N^k) \quad (64)$$

where we have assumed that $j < i$. Hence we can write (63) as

$$W_i^{k+1} = \max\{f_i(\mathbf{W}^{k,i-}), R_i\} \quad (65)$$

Note the policy used at each step π_{k+1} is itself varying with time, but is only a function of the neighbor nodes of i at k . Also note that though in general f_i is a function of all nodes in the system other than i , often only some subset of these nodes are actually neighbors of i and affect the update computation.

Lemma 5.7: For every i , $f_i(\mathbf{W}^{k,i-})$ is component-wise continuous and piecewise linear, of non-negative monotonically non-decreasing slope.

Proof: The proof of Lemma 5.7 is in Appendix E. ■

We pick an arbitrary node from which to observe the update of all the nodes in the r-r system. Without loss of generality, we consider node N . Let \mathbf{W}^n be the node values after the n th computation and transmission of W_N^n (the components of \mathbf{W}^n are as in (63)).

Definition 5.5: Define the mapping $T: \mathbb{R}^N \rightarrow \mathbb{R}^N$ as

$$\mathbf{W}^{n+1} = T(\mathbf{W}^n)$$

where

$$\begin{aligned} W_1^{n+1} &= \max\{f_1(W_2^n, \dots, W_N^n), R_1\} \\ W_2^{n+1} &= \max\{f_2(W_1^{n+1}, W_3^n, \dots, W_N^n), R_2\} \\ &\vdots \\ W_{N-1}^{n+1} &= \max\{f_{N-1}(W_1^{n+1}, \dots, W_{N-2}^{n+1}, W_N^n), R_{N-1}\} \\ W_N^{n+1} &= \max\{f_N(W_1^{n+1}, \dots, W_{N-1}^{n+1}), R_N\} \end{aligned} \quad (66)$$

Note that T is fixed for all n . The fact that this T mapping is fixed is the reason we have introduced this notion of a r-r system. It allows for a fixed-point argument which we now develop. Our procedure is to demonstrate the continuity of the mapping T , show that \mathbf{V}^π is a unique fixed point of T , and then use continuity of T to show that the limit point of \mathbf{W}^n is a fixed point of T . We can then conclude that this limit point is the value function.

Lemma 5.8: The mapping T is component-wise continuous for each output.

Proof: The assertion of Lemma 5.8 follows from (66), Lemma 5.7, the fact that $\max\{\cdot, \cdot\}$ is continuous in its arguments, and the fact that composition of continuous functions results in a continuous function. ■

Lemma 5.9: $\mathbf{V}^\pi = T(\mathbf{V}^\pi)$, and there are no other fixed points of T .

Proof: The proof of Lemma 5.9 is in Appendix F. ■

Lemma 5.10: $W_i^* := \lim_{k \rightarrow \infty} W_i^k = V_i^\pi, \forall i$

Proof: We have

$$\mathbf{W}^* = \lim_{k \rightarrow \infty} \mathbf{W}^{k+1} = \lim_{k \rightarrow \infty} T(\mathbf{W}^k) = T(\lim_{k \rightarrow \infty} \mathbf{W}^k) = T(\mathbf{W}^*)$$

where the third equality holds because of Lemma 5.7, Lemma 5.8, and Theorem 4.10 of [30]. Hence any limit point of \mathbf{W}^k must be a fixed point of T . But by Lemma 5.9, the only fixed point of T is \mathbf{V}^π , and the result follows. ■

Based on Lemma 5.2-Lemma 5.10 we prove Theorem 5.1.

Proof of Theorem 5.1

Proof of 1. By Lemma 5.3 and Lemma 5.4, after finite time n_p we have

$$V_i^n \geq V_i^\pi, \quad \forall n \geq n_p, \forall i \in \Omega \quad (67)$$

Combining Lemma 5.6 and Lemma 5.10, we have

$$\lim_{n \rightarrow \infty} \bar{D}_i^{v_n^i} \leq \lim_{n \rightarrow \infty} W_i^{u_n^i} = \lim_{k \rightarrow \infty} W_i^k = V_i^\pi, \quad \forall i \in \Omega \quad (68)$$

Lemma 5.3 (1) and (68) ensure that

$$\lim_{n \rightarrow \infty} V_i^n \leq \lim_{n \rightarrow \infty} \bar{D}_i^{v_n^i} \leq V_i^\pi, \quad \forall i \in \Omega \quad (69)$$

Inequalities (67) and (69) imply that $\lim_{n \rightarrow \infty} V_i^n = V_i^\pi$.

Proof of 2. Because of (51) there exists n_p after which the node values at each node i and its neighbors $\mathcal{N}(i)$ are in the order of some local index policy that coincides with the order of an optimal index policy. At each event $n \geq n_p$ this local

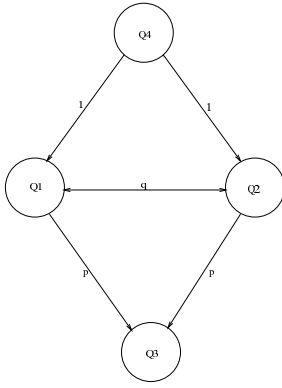


Fig. 2. System with optimal distributed policy

ordering can change, but the new local order coincides with the order of an optimal index policy. Hence, at each node i an optimal action is taken at each event $n \geq n_p$. Consequently, the distributed policy is optimal.

Proof of 3. Because of (51), (52), and Lemma 5.3, there exists n_p after which node values for each node i and its neighbors $\mathcal{N}(i)$ are in the order of the unique optimal local index policy and there are no subsequent changes in this order. After n_p , once a full round of node updates occurs, each node has its correct V_i^π and its correct local index policy, which subsequently do not change. \square

We give an example where Algorithm 2 does not converge to an optimal index policy, but still converges to an optimal distributed policy.

Example

Consider the system of Figure 2 with parameters $0 < p < 1$, $0 < q \leq 1$, $c_1 = c_2 = c_4 = 1$, and $R_1 = R_2 = R_4 = 0$, $R_3 = R$. Assume node recomputations and successful value transmissions occur in the order $[1, 4, 2, 4, 1, 4, 2, 4, \dots]$. Then node 1 and 2 updates are still represented by (60). The policy at node 1 is fixed at $\pi_t^1 = (3, 2, 1), \forall t$, and the policy at node 2 is fixed at $\pi_t^2 = (3, 1, 2), \forall t$. But at each update of node 4, the policy changes, as the ranking of the values of nodes 1 and 2 alternate. That is, the policy computed at node 4 is

$$\pi_t^4 = \{(1, 2, 4) (2, 1, 4) (1, 2, 4) (2, 1, 4) (1, 2, 4) \dots\} \quad (70)$$

Of course, $\pi_t^3 = r_3$. The overall policy is $\Pi_t = (\pi_t^1 \pi_t^2 \pi_t^3 \pi_t^4)$, which is not a stationary policy, due to π_t^4 . It is, however, an optimal distributed policy.

VI. DISTRIBUTED DYNAMIC PROGRAMMING FORMULATION

We develop a distributed algorithm, different from but related to Algorithm 2, using the methodology of Distributed Dynamic Programming (DDP) [3]. DDP is a technique for solving dynamic programming problems using distributed computation. The technique may only be used for problems formulated so that a standard dynamic programming equation applies, and for which a suitable partitioning of the state space among processors (“computation centers”) can be made. See

[3] for a description of DDP, and [17] for further details on applying DDP to Problem (\mathbf{P}_1) .

A. Solution of Problem (\mathbf{P}_1) Using DDP

The model of Problem (\mathbf{P}_1) is a standard controlled Markov chain with finite state space 2^Ω and action space $i \in \Omega$. Dynamic programming can be directly applied to Problem (\mathbf{P}_1) on state space 2^Ω , but this approach is inefficient and does not lead to a direct application of DDP. This is because it is not possible to define useful computation centers as in [3] through partition of $S = 2^\Omega$ (on this state space, to define computation centers in a way that leads to the application of DDP, we would need to know *a priori* the optimal priority list of nodes as dictated by Theorem 2.1). Our approach is to use the index structure of an optimal policy demonstrated in Theorem 2.1 to define a new state space on which DDP can be applied.

To solve Problem (\mathbf{P}_1) using DDP, we proceed in three steps.

- 1) Formulate a new problem, Problem (\mathbf{F}_1) below
- 2) Show that an optimal policy for Problem (\mathbf{F}_1) can be mapped to an optimal distributed policy for Problem (\mathbf{P}_1)
- 3) Apply DDP to Problem (\mathbf{F}_1)

Below we present the details of each of the steps above.

1) *Formulation of Problem (\mathbf{F}_1)* : We formulate Problem (\mathbf{F}_1) .

Problem (\mathbf{F}_1)

The state space is Ω , the set of nodes. The policy space is the set of all local index policies π^i for all $i \in \Omega$. When transmitting in state i , a cost c_i is incurred and transition to a new state $j \in \Omega$ occurs. Define $P_{ij}^{\pi^i}$ to be the probability of transition from state i to state j , $j \in \mathcal{N}(i)$, under policy π^i . Then,

$$P_{ij}^{\pi^i} = \sum_{S: \pi^i(S)=j, r_j} P^i(S|i) \quad (71)$$

When we retire in state i , the process terminates and a reward R_i is received. The objective is to choose for each state i the local index policy π^i to maximize

$$E \left\{ R_{i(\tau)} - \sum_{t=1}^{\tau-1} c_{i(t)} \right\} \quad (72)$$

where τ is the time when the transmission process is terminated, and $i(t)$ is the state at time t .

2) *Mapping of Optimal Policy from Problem (\mathbf{F}_1) to Problem (\mathbf{P}_1)* : We define a mapping of policies for Problem (\mathbf{F}_1) to policies for Problem (\mathbf{P}_1) .

Mapping 1: A general (possibly time-varying) policy for Problem (\mathbf{F}_1) consists of a sequence of local index policies for each $i \in \Omega$. For each t we define

$$\Pi_t = \{\pi_t^1 \pi_t^2 \dots \pi_t^N\} \quad (73)$$

A policy for Problem (\mathbf{F}_1) is then specified as

$$\Pi = \{\Pi_1 \Pi_2 \dots \Pi_t \dots\} \quad (74)$$

We map a policy for Problem (\mathbf{F}_1) represented as (74) into a distributed index policy for Problem (\mathbf{P}_1) (cf. Definition 4.1) in the obvious way. That is, at each time t the local index policy π_t^i at each node $i \in \Omega$ is the same for Problem (\mathbf{P}_1) as for Problem (\mathbf{F}_1) .

We show that mapping an optimal policy of Problem (\mathbf{F}_1) by Mapping 1 leads to an optimal policy for Problem (\mathbf{P}_1) .

Lemma 6.1: If Π is an optimal policy for Problem (\mathbf{F}_1) , then Π mapped to a policy for Problem (\mathbf{P}_1) using Mapping 1 is an optimal distributed policy for Problem (\mathbf{P}_1) .

Proof: The assertion of Lemma 6.1 follows from Mapping 1 and the fact that an optimal policy for Problem (\mathbf{P}_1) is of the index type. ■

3) *A Distributed Dynamic Programming Implementation for Problem (\mathbf{F}_1) :* We formulate a DDP solution of Problem (\mathbf{F}_1) by translating the notation of [3] into our notation. Each node i in our model is a computation center for itself alone, so that $S_i = i, \forall i \in \Omega$. We associate

$$\begin{aligned} S &= \Omega \\ C &= \{\text{Set of all priority orderings of } \mathcal{N}(i), \forall i \in \Omega\} \\ U(i) &= \{\text{Set of all priority orderings of } \mathcal{N}(i)\} \\ F &= \{-\infty, R_{max}\} \end{aligned}$$

We also note the notational correspondences $V \leftrightarrow J$, and $V_i^n \leftrightarrow J_i^t$, which are equivalent assuming event n occurs at time t .

A neighbor of a node in the sense of [3] corresponds to our notion of neighbor. That is, the neighbors of node i are those nodes with positive marginal probability of receiving the packet when i transmits. These are precisely the nodes whose values can affect node i 's update. However, the notation for the neighbors of i used in [3], $N(i)$, does not include i , whereas for our notation $i \in \mathcal{N}(i)$.

For Problem (\mathbf{F}_1) we define the H function of [3] $\forall i \in \Omega, \pi^i \in U(i)$ as

$$H(i, \pi^i, V) = -\max \left\{ -c_i + \sum_{S \subseteq \mathcal{N}(i)} P(S|i) V(\pi^i(S)), R_i \right\} \quad (75)$$

The negative sign in the RHS of (75) is used to conform to the convention of [3] that the goal is to minimize cost. Thus, the dynamic programming equation for Problem (\mathbf{F}_1) is

$$V_i^\pi = \min_{\pi^i} H(i, \pi^i, V^\pi) \quad (76)$$

Note that H is monotone in V in the sense of [3].

The update for state $i \in \Omega$ by DDP at event n is

$$V_i^n = \min_{\pi^i} H(i, \pi^i, V^n) \quad (77)$$

We suppose that Assumption 1 of [3] is true for node updates and transmissions. The functions

$$\underline{V}_i := -R_{max}, \quad \forall i \in \Omega \quad (78)$$

$$\bar{V}_i := 0 \quad (79)$$

satisfy part (i) of Assumption 2 of [3] due to (75) and (77), and part (ii) of Assumption 2 due to (76) and the standard result on value iteration for a dynamic program. We require that, when running DDP for Problem (\mathbf{F}_1) , we start with a value between \underline{V} and \bar{V} at each node.

This completes the DDP formulation of Problem (\mathbf{F}_1) . We now briefly state the results from [3] which apply to Problem

(\mathbf{F}_1) . Assumption 1, Assumption 2, and the requirements of Proposition 1 in [3] are satisfied. Proposition 1 of [3], translating into our notation, then implies that

$$\lim_{n \rightarrow \infty} V_i^n = V_i^\pi, \quad \forall i \in \Omega \quad (80)$$

Because for fixed π^i update (75) is continuous in the components of V , the requirements of Proposition 3 of [3] are also satisfied. Translating to our notation, we obtain from Proposition 3 that there exists a $\bar{n}_i > 0$ such that for all $n \geq \bar{n}_i$, if a local index policy π_n^i satisfies

$$H(i, \pi_n^i, V_i^n) = \min_{\pi^i} H(i, \pi^i, V_i^n), \quad \forall i \in \Omega \quad (81)$$

then

$$H(i, \pi_n^i, V^\pi) = \min_{\pi^i} H(i, \pi^i, V^\pi) \quad (82)$$

Equation (82) states that π_n^i takes an optimal action for state i at event n .

Thus, we have shown how Problem (\mathbf{F}_1) can be solved with the DDP methodology of [3] using update (77).

B. Relation Between DDP Solution to Problem (\mathbf{F}_1) and Algorithm 2

Event 1 of Algorithm 2 corresponds precisely to the transmit state, which is update 1) of [3]. Event 2 of Algorithm 2 is related to update 2) of [3]. However, the update equations (48) for Algorithm 2 and (77) for DDP are not the same. For an update at node $i \in \Omega$ at event $n+1$, the key difference is that in (77) the value V_i^n affects the updated value V_i^{n+1} , whereas in (48) it does not. Another difference, relatively minor, between the DDP formulation and Algorithm 2 is that where Algorithm 2 assumes Events 1 and 2 occur infinitely often, the DDP formulation of Problem (\mathbf{F}_1) requires the somewhat more restrictive Assumption 1.

Though Algorithm 2 and the DDP formulation differ in the ways just mentioned, the results proved for each algorithm are quite similar. As remarked in (80), Proposition 1 in [3] implies that

$$\lim_{n \rightarrow \infty} V_i^n = V_i^\pi, \quad \forall i \in \Omega \quad (83)$$

Equation (83) shows asymptotic convergence to the value function for Problem (\mathbf{F}_1) , and hence for Problem (\mathbf{P}_1) , and is similar to result Theorem 5.1, (1.). The result of Proposition 3 (82) is similar to our Theorem 5.1, (2.). No results similar to Theorem 5.1, (3.) for Problem (\mathbf{P}_1) under the DDP formulation follow directly from a result of [3].

As these results show, once Theorem 2.1 is proved for Problem (\mathbf{P}_1) , application of Distributed Dynamic Programming on an appropriate state space provides a new distributed algorithm for Problem (\mathbf{P}_1) with certain properties nearly equivalent to those of Algorithm 2.

VII. DISTRIBUTED RANK METHOD

We define a third distributed algorithm for Problem (\mathbf{P}_1) , which is similar in spirit to Algorithm 2, except that a different method is used to update the node value function. We define the following

Algorithm 3: At each event time, any number of the following two events can occur.

Event 1 A node i receives V_j from a neighbor $j, j \in \mathcal{N}(i)$.

Event 2 A node i recomputes V_i using the current $V_{i,j}$ values, as follows.

- 1) The set of $V_{i,j}$ and V_i values are ranked, high-to-low. A local index policy $\hat{\pi}$ for i is created which uses this ranking (ties are broken arbitrarily). If $V_i \geq V_{i,j} \forall j \in \mathcal{N}(i)$, then $\hat{\pi}$ uses *any* ranking of the neighbors, so long as i is given a ranking below at least one neighbor.
- 2) The following is computed.

$$V_i^n = \max \left\{ \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \neq i} P^i(S|i) V_{i, \hat{\pi}(S)}^n}{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \neq i} P^i(S|i)}, R_i \right\} \quad (84)$$

We assume events 1 and 2 occur infinitely often at each node i .

Before we proceed with the analysis of the convergence properties of Algorithm 3, we define certain types of events associated with the execution of Algorithm 3.

Definition 7.1: We call a *node reset* the event where V_i is equal to or larger than all neighbors in Step (1) of Algorithm 3, at which point the algorithm sets i 's rank somewhere lower than the highest neighbor, and the left-hand term using $\hat{\pi}$ is the larger term in (84) in Step (2).

When a node reset occurs for node i , the new ranking used for $\hat{\pi}$ is an implementation decision. All of the following results for Algorithm 3 are valid for any choice of $\hat{\pi}$ when resetting, so long as node i is not ranked highest. Of course, the choice of $\hat{\pi}$ affects the convergence properties of Algorithm 3. An example of this choice would be to always make i the node of lowest rank among neighbors when resetting. Another choice would be to always make i the second highest ranked node among the neighbors when resetting.

Definition 7.2: When a node value recomputation using (84) occurs at node i which is *not* a node reset, we say a *standard update* has occurred.

Definition 7.3: When a standard update occurs without the retirement value R_i being the optimal choice, so that the left term in the RHS of (84) is chosen, we say a *rank update* has occurred.

We state our main results for Algorithm 3 in the following theorem.

Theorem 7.1: For Algorithm 3 with any initial state s.t. $0 \leq V_i^0 \leq R_{\max}$ and $0 \leq V_{i,j}^0 \leq R_{\max}$ for all $i, j \in \Omega$, we have

- 1) $\lim_{n \rightarrow \infty} V_i^n = V_i^\pi, \forall i \in \Omega$
- 2) There exists an event $n_p < \infty$ s.t. $\Pi = \{\Pi_{n_p}, \Pi_{n_p+1}, \Pi_{n_p+2}, \dots\}$ is an optimal distributed policy
- 3) If

$$V_i^\pi \neq V_j^\pi, \forall i \in \Omega, j \in \mathcal{N}(i) \quad (85)$$

then there exists an event $n_p < \infty$ s.t.

- a) $V_i^n = V_i^\pi, \forall n \geq n_p, \forall i \in \Omega$
- b) $\Pi = \{\Pi_{n_p}, \Pi_{n_p+1}, \Pi_{n_p+2}, \dots\}$ has an associated index policy that is optimal

Before proving Theorem 7.1, we present a series of lemmas which will be used in the proof. But first we note that the monotonicity property in the sense of Lemma 5.2 does not

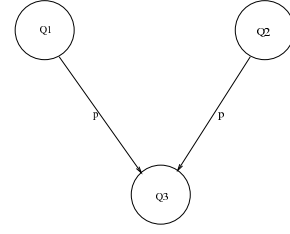


Fig. 3. System Demonstrating Violation of Monotonicity for Rank Update

hold in general for the update of Algorithm 3. We show this by the following example.

Example

Consider the system of Figure 3. Given $R_3, V_3^0, V_{3,1}^0$, and $V_{3,2}^0$, assume at time 1 there is a node update using (84) for node 3. We consider the following two cases. The values used in the update are: Case 1) $R_3, V_3^0, V_{3,1}^0$, and $\hat{V}_{3,2}^0$, and Case 2) $R_3, V_3^0, V_{3,1}^0$, and $\tilde{V}_{3,2}^0$. Assume $R_3 = 0$ and that

$$\tilde{V}_{3,2}^0 < V_3^0 < \hat{V}_{3,2}^0 < \frac{-c_3}{p} + V_{3,1}^0 \quad (86)$$

Case 1 From (84) we have $\pi_3^1 = (1, 2, 3)$. Then $\hat{V}_3^1 = \frac{-c_3 + pV_{3,1}^0 + (1-p)p\hat{V}_{3,2}^0}{p + (1-p)p}$.

Case 2 From (84) we have $\pi_3^1 = (1, 3, 2)$. Then $\tilde{V}_3^1 = \frac{-c_3 + pV_{3,1}^0}{p}$.

Noting that $\hat{V}_{3,2}^0 < \frac{-c_3}{p} + V_{3,1}^0$, we find that $\hat{V}_3^1 < \tilde{V}_3^1$. But since $\tilde{V}_{3,2}^0 > \hat{V}_{3,2}^0$ and all else is equal in the two cases, update monotonicity is violated.

To make the example concrete, let $V_3^0 = 5, V_{3,1}^0 = 10, p = .5, c_3 = 1$. Let $\tilde{V}_{3,2}^0 = 3$, so $\tilde{V}_3^1 = -2 + 10 = 8$. Let $\hat{V}_{3,2}^0 = 6$, so $\hat{V}_3^1 = \frac{-1 + (.5)10 + .25(6)}{.5(1.5)} = 7\frac{1}{3} < \tilde{V}_3^1$, violating monotonicity. \square

Because the update monotonicity property does not hold for Algorithm 3, we must take a very different approach to proving its convergence than what we used for Algorithm 2. In the following lemma we prove what may appear as a technical result of unclear utility, but which proves crucial in the approach we take. It says that when two index policies share the top of a priority list, but differ further down, then their updated values can be usefully related.

Lemma 7.1: Assume we are given two local non-retiring index policies $\hat{\pi}$ and $\tilde{\pi}$ and a full set of neighbor values $\{V_{i,j}\}$ for node i . Let $A \subset \mathcal{N}(i)$ be the nodes which $\hat{\pi}$ ranks above node i . Assume that policy $\tilde{\pi}$ uses this same ranking for the nodes in A , but that there is also a set $B \subset \mathcal{N}(i) - A, B \neq \emptyset$, each member of which $\tilde{\pi}$ ranks above node i and below the nodes of A .

Define $Y_1 := \max_{j \in B} V_{i,j}$, and $Y_2 := \min_{j \in B} V_{i,j}$. Then there exists $0 \leq a < 1$ such that

$$(1-a)V_i^{\hat{\pi}} + aY_2 \leq V_i^{\tilde{\pi}} \leq (1-a)V_i^{\hat{\pi}} + aY_1 \quad (87)$$

Proof: The proof of Lemma 7.1 is in Appendix G. \blacksquare

In the next lemma, we prove a straightforward fact about a rank update, which is that the updated value is upper bounded by the highest neighbor value minus the cost c_i .

Lemma 7.2: Suppose there is a rank update for node i at event n , with neighbor node values $V_{i,j}^{n-1}, j \in \mathcal{N}(i)$, and neighbor ranking $\tilde{\pi}$. Define $V_{max} = \max_{j \in \mathcal{N}(i): j >_{\tilde{\pi}} i} V_{i,j}^{n-1}$. Then

$$V_i^n \leq V_{max} - c_i \quad (88)$$

Proof: The proof of Lemma 7.1 is in Appendix H. ■

As before, we use $V^\pi(\cdot)$ to denote the optimal value function for Problem (P₁).

In the following lemma, we use a continuity property of the Algorithm 3 update to translate bounds on neighbor values into bounds on the updated value.

Lemma 7.3: Suppose π_j is an optimal local policy for node j , and let $B := \{k \in \mathcal{N}(j) : k >_{\pi_j} j\}$. Then for any $\epsilon > 0$ there is $\delta > 0$ such that

$$V_{j,k} \geq V_k^\pi - \delta, \forall k \in B \implies V_j^{\pi_j} \geq V_j^\pi - \epsilon \quad (89)$$

Suppose $\tilde{\pi}_j$ is a local policy (possibly suboptimal) for node j , let $\tilde{B} := \{k \in \mathcal{N}(j) : k >_{\tilde{\pi}_j} j\}$, and assume that $\tilde{B} \subset B$. Then for any $\epsilon > 0$ there is $\delta > 0$ such that

$$V_{j,k} \leq V_k^\pi + \delta, \forall k \in \tilde{B} \implies V_j^{\tilde{\pi}_j} \leq V_j^\pi + \epsilon \quad (90)$$

Proof: The proof of Lemma 7.3 is in Appendix I. ■

The fact that Lemma 7.1 guarantees the existence of an $0 \leq a < 1$ satisfying (87) is key to the arguments we develop in this section. We simplify the use of this fact by defining a maximum over all such a 's.

Definition 7.4: For a given node $i \in \Omega$ and a pair of local index policies $\tilde{\pi}$ and $\hat{\pi}$, let $a(i, \tilde{\pi}, \hat{\pi})$ be given by (120), (121), and (126). Define

$$a_{max} := \max_{i \in \Omega, \tilde{\pi}, \hat{\pi}} a(i, \tilde{\pi}, \hat{\pi}) \quad (91)$$

Note that $0 \leq a_{max} < 1$.

The following lemma is our main substantive result for the rank algorithm, and leads directly to Theorem 7.1.

Lemma 7.4: For Algorithm 3 we have $\lim_{n \rightarrow \infty} V_i^n = V_i^\pi, \forall i \in \Omega$

Proof: The proof of Lemma 7.4 is in Appendix J. ■

We are now ready to present the proof of Theorem 7.1.

Proof: [of Theorem 7.1]

Proof of 1. Relation 1. follows directly from Lemma 7.4.

Proof of 2. Because of 1. there exists n_p after which the node values at each node i and its $\mathcal{N}(i)$ are in the order of some local index policy with an optimal associated index policy. At each event $n \geq n_p$ this local ordering can change, but it always corresponds to some optimal associated index policy. Hence, at each node i an optimal action is taken at each event $n \geq n_p$. Hence, the distributed policy is optimal.

Proof of 3. Because of 1. and Lemma 7.4, there exists n_p after which node values for each node i and its $\mathcal{N}(i)$ are in the order of the unique optimal local index policy. and there are no subsequent changes in this order. Subsequent to n_p , once a full round of node updates occurs, each node has its correct V_i^π and its correct local index policy, which subsequently do not change. ■

Discussion

It is interesting to note the relation of Algorithm 3 to the results in [8]. Major differences include the on-going nature of the node updates in Algorithm 3, and that node values are

Alg.	V_i^n Update
Alg. 2	$\max \left\{ \max_{\pi} \left\{ \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i} P^i(S i) V_{i,\pi(S)}^{n-1}}{\sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i} P^i(S i)} \right\}, R_i \right\}$
DDP	$\max \left\{ \max_{\pi} \left\{ -c_i + \sum_{S \subseteq \mathcal{N}(i)} P^i(S i) V_{i,\pi(S)}^{n-1} \right\}, R_i \right\}$
Alg. 3 (Rank)	π set to rank order of $V_{i,j}^{n-1}$ $\max \left\{ \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i} P^i(S i) V_{i,\pi(S)}^{n-1}}{\sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i} P^i(S i)}, R_i \right\}$

TABLE I
UPDATE FUNCTION FOR THREE DISTRIBUTED ALGORITHMS

interpreted as the expected reward at a given node, whose estimates are constantly being updated by (84). The direction flipping action of [8] corresponds to resetting in Algorithm 3. The proof of Theorem 7.1 is more difficult than the analogous result in [8], due to the fact that in Algorithm 3 node values can change not just when being reset but also when normally computing updated estimates.

VIII. ALGORITHM COMPARISON AND SIMULATION RESULTS

For ease of reference, we summarize the update functions of this paper's three algorithms in Table I. All three algorithms run in a similar fashion, with the main difference being this update function. In terms of computational complexity of the update, Alg. 2 and the DDP algorithm are roughly comparable, while the Rank algorithm is somewhat less complex. In Theorem 5.1, the DDP results, and Theorem 7.1 we have shown finite-time convergence of each algorithm to an optimal policy in a static network. There remain important questions about how each algorithm dynamically behaves during convergence, such as the convergence rate, which have not been addressed in this paper. Understanding dynamic behavior would allow us to characterize these algorithms more completely.

An ad hoc network simulation has been conducted to compare performance of the three distributed algorithms of Table I, along with the performance of other algorithms in the literature. The simulation is setup so that the frequency of packet arrival justifies a route-maintenance approach. The density of mobiles in the simulated network is chosen so that the local broadcast nature of transmission is beneficial. Hence, the results presented here are not conclusive in comparing the different approaches, but they do represent an important class of networks for which our approach performs very well.

The network consists of 12 mobiles on a field 2 km x 2 km. Each mobile moves at the same constant speed, which can be varied over different simulation runs. Each mobile chooses a waypoint which is a random point in the field of operation, moves toward this waypoint at a fixed speed, and then chooses a new random waypoint once the previous one is reached. Messages arrive at a fixed rate at each mobile, and they all have the same destination node, which is also

moving. Mobiles transmit each packet as it arrives at fixed power, and take statistics on which other mobiles receive the packet by keeping a moving average of the success rate to each other mobile. Mobiles are assumed to be transmitting pilot signals from which the active set of neighbors can be determined for both control signaling and decoding attempts. Transmission instances are assumed synchronous across all mobiles. Each mobile has a single omnidirectional antenna for both transmitting and receiving.

Transmission channels and interference are modelled using the method from the 3GPP2 1xEV-DO Evaluation Methodology [1] for CDMA transmission. In the present simulation, only Channel A (a 3kmph single-path Rayleigh fading channel) is used, and short-term curves (PER vs. E_b/N_0) generated for this channel are used to determine packet success. When variation in velocity is indicated in the simulation results, this indicates the rate of network topology change, not of the fading process. Signal propagation loss is determined using a power of 3.5 dropoff in distance. Path loss to each neighboring mobile consists of a propagation loss, time-varying shadowing, and short-term fading. The propagation loss is based on distance, the time-varying shadowing is assumed .5 correlated across neighbors, and the short-term fading is assumed independent across neighbors. Each time instant, the power received at each node is determined and an interference term is used in the E_b/N_0 computation for each mobile packet transmission. We model a carrier frequency of 1.9GHz, a chip rate of 1.25Mcps, and a spreading gain of 20, which implies an approximate 60kbps transmission rate. We assume each packet is 50msec in duration.

Each transmission is assumed to cost a fixed energy of 1 unit and occur at a fixed power. After each transmission from a mobile, the routing protocol uses the set of neighbor mobiles successfully receiving the packet to determine whether to retransmit the packet or to pass control of the packet to a neighbor mobile. The energy cost of the packet reception and control signaling is assumed negligible in these simulations.

As the mobiles move toward their waypoints, and as shadowing and fading vary with time, transmission success probability to each neighbor dynamically changes. Transmission probabilities are estimated for all algorithms using a moving window average of recent success frequency. These successes are kept individually for each neighbor and assumed independent. Though this is not exactly true physically (nor in the simulation transmission model itself), the effect of fading over packet transmission times often means this model is a reasonable approximation and useful for making routing decisions. Note that nothing in the optimality of these algorithms depends on such independence, but it is assumed for computational simplicity of the routing algorithms used in this simulation.

We simulate the performance of six algorithms. The first three are those presented in this paper and summarized in Table I. The remaining three are version of well-known algorithms: Distributed Bellman-Ford (DBF) [4] with a hop-count metric (DBF-HC), DBF with an expected energy cost metric (DBF-EN), and the method of Gafni-Bertsekas (GB) [8], which has also been adapted for use in the TORA algorithm

[22]. In DBF-HC, the path with fewest links to the destination is used. In DBF-EN, a metric for each link is determined based on the expected energy to transmit across the link, which in our model is the inverse of the success probability, and the path of smallest cumulative expected energy is used. In GB, the same path is used until a link in the path is broken, at which time a distributed algorithm runs which is guaranteed to determine a new path to the destination. In the DBF-HC and GB algorithms, a strict binary decision is made as to whether or not each link is connected. In our simulation, we use a fixed probability threshold to make this determination.

We simulate two cases, that of a *loaded* network, and that of a *partially loaded* network. In the case of a loaded network, it is assumed that each mobile is transmitting and generating interference nearly all of the time, even when the mobile does not currently have any packet destined for the destination being simulated. This accounts for the fact that though we simulate packets for just one destination, in the actual network there may be many such destinations and their packets are constantly flowing in the network. The loaded network is the case where transmission opportunities are close to saturation. In the partially loaded network, we assume that each mobile transmits only a fraction of the time, with these times chosen randomly, when generating neighbor interference.

In summary, the algorithms simulated were:

- Algorithm 2
- Algorithm 3, Rank Algorithm
- DDP Algorithm
- Distributed Bellman-Ford (DBF-HC), where link cost is always 1 (i.e. count number of hops)
- Distributed Bellman-Ford (DBF-EN), where link cost is expected energy cost to transmit across the link
- Gafni-Bertsekas algorithm (GB), related to the TORA class of ad hoc algorithms

We first show the nature of convergence of the average energy expended per packet (EPP) to reach its destination. Because each packet transmission requires one unit of energy, average EPP also indicates the number of times on average a transmission attempt was made per packet before it finally reaches the destination. A running total of this average is recorded in the simulation each time step. The average EPP vs. time value is plotted in Figure 4 for the 3kmph case for all six algorithms for a loaded network. The overhead signaling for each of these algorithms was not accounted for (e.g. for GB upon occasion the entire path must be updated with extensive control signaling, but we ignore the cost of that here). Because the energy cost of a packet is only recorded once it reaches the destination, early measurements of EPP are biased toward good packets, and hence are optimistic. As the simulation progresses, the average settles to a steady-state value indicating the correct overall EPP.

The three algorithms using stochastic routing outperform the three algorithms that do not by nearly 3dB. We note in passing that this is also the order of improvement claimed for soft handoff in cellular systems [32]. Interestingly, the Alg2, DDP, and Rank algorithms all converge to the same value, though in this example the DDP algorithm is slightly less efficient while doing so. It is important to recognize that,

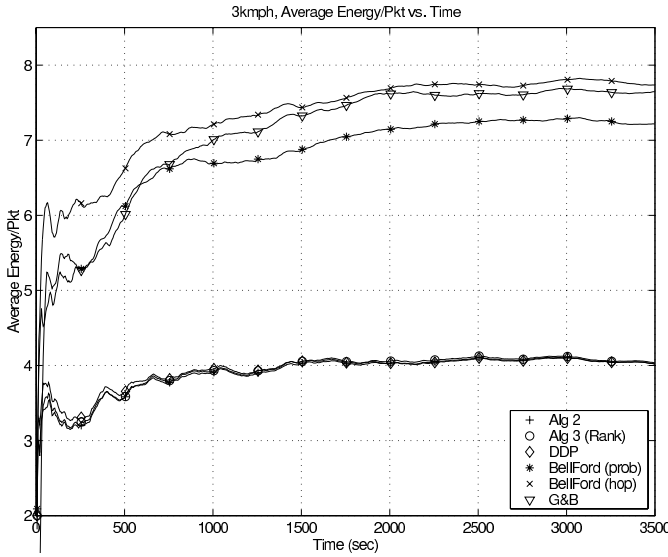


Fig. 4. Convergence of Average Energy per Packet for Six Distributed Routing Algorithms, Loaded Network

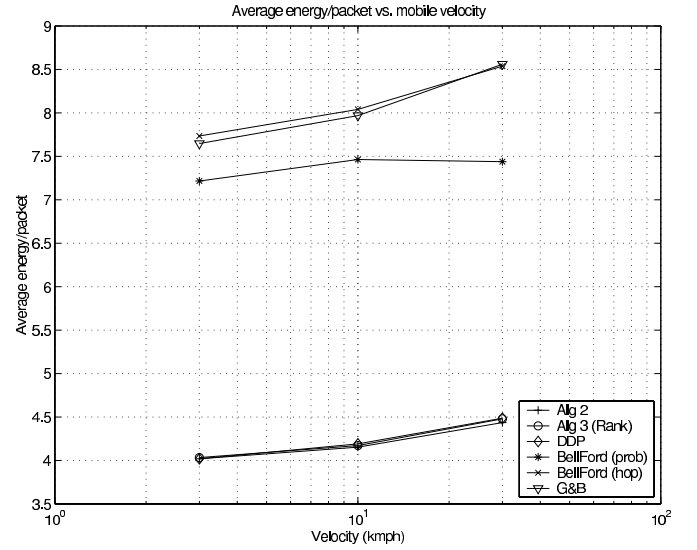


Fig. 5. Average Energy per Packet vs. Velocity for Six Distributed Routing Algorithms, Loaded Network

though we have proved that these three algorithms converge to the optimal routing algorithm in the steady-state, their dynamic performance is potentially different in this more realistic non-stationary context with accurate channel models. As expected, DBF-EN moderately outperforms DBF-HC, while GB converges to a value slightly better than DBF-HC. It is difficult to predict in general how DBF-HC and GB will compare to each other. There is no guarantee that the hop-count metric is actually better than the path in use by GB, because, though the GB path is chosen in a somewhat arbitrary manner, the actual cost of the hop-count path may also vary widely based on the link channel quality. The DBF-EN algorithm takes the link quality into account, and so should be the best of the algorithms which update a fixed path.

Figure 5 shows average EPP for the six algorithms as a function of mobile velocity, which is used to update network topology. Note that in general the performance of each algorithm suffers as the velocity increases. The basic effect is that the estimate of transmission probability from past events becomes less relevant to future transmission success as the topology changes more and more rapidly. For the three stochastic routing algorithms and for DBF-EN, the inaccuracy in probability estimation translates directly into routing inefficiency. For DBF-HC and GB, which use a channel quality threshold, the effect is to postpone when link state change detection occurs, and the quality of the path in use suffers accordingly. Note also that Alg 2 very slightly outperforms DDP and Rank for this case at higher velocity.

Figure 6 also shows average EPP for the six algorithms as a function of mobile velocity, but for the case of a partially loaded network. Note that the average EPP is about half that of Figure 5, but the ratio of each algorithms performance is roughly comparable. This indicates that the advantages of stochastic routing extend to the case of multiple loading levels of a network. One question to pursue about this plot is why the DBF-HC, DBF-EN, and GB algorithms actually improve at

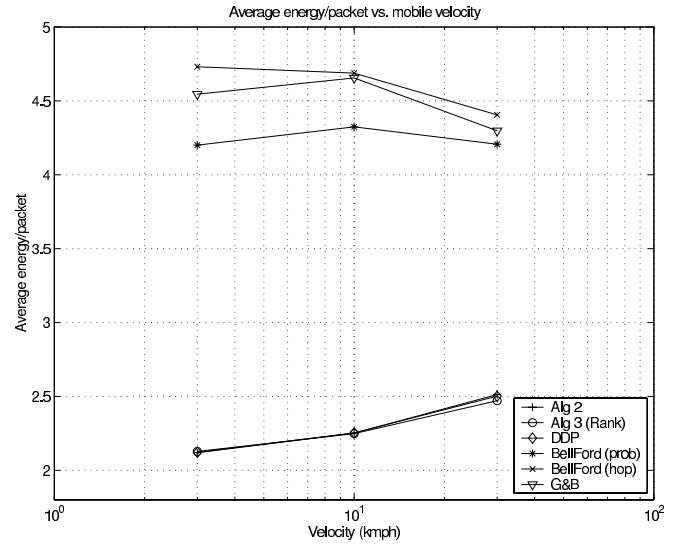


Fig. 6. Average Energy per Packet for Six Distributed Routing Algorithms, Partially Loaded Network

the highest velocity. Along these lines note that though higher velocity has the effect of impeding probability estimation, it can also be beneficial in keeping the network sufficiently mixing to keep fixed paths from clogging up.

Finally, we consider delivery delay per packet. Delivery delay is different from EPP because packets can also incur queuing delay at each mobile. This queuing delay is generally worse for the fixed-path algorithms, because packets locally tend to take the same path. Note that if load-balancing were implemented to improve delay, then EPP would generally increase, at least for DBF-EN. In contrast, the stochastic routing approach naturally leads to more spreading out of packet transmissions among mobiles, greatly reducing the need for an explicit load balancing algorithm.

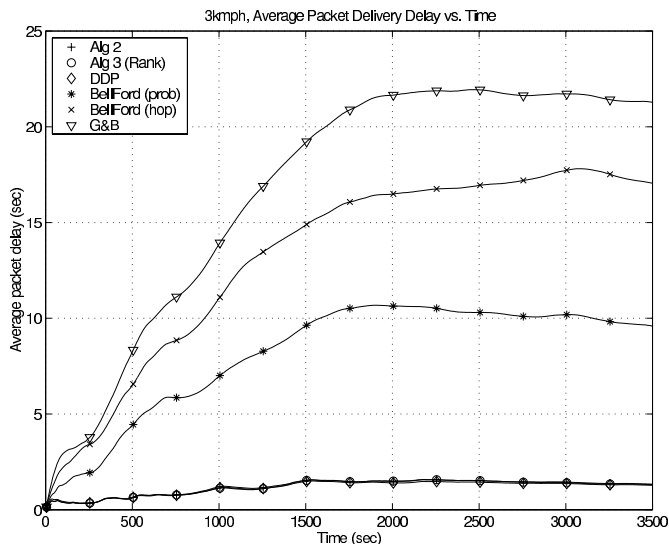


Fig. 7. Convergence of Delivery Delay per Packet for Six Distributed Routing Algorithms, Loaded Network

Figure 7 shows the time trace of average delivery delay per packet for each algorithm in the loaded network for the case of 3 kmph. Note that the stochastic routing delay performance is substantially better than that of the other algorithms. Also, DBF-EN outperforms DBF-HC and GB more significantly in delay. This is because the expected energy metric leads to more variation in the path used, which leads to less congestion. Though none of these algorithms has been optimized for delay, the delay statistic is still of interest in understanding algorithm performance. It is useful when an algorithm achieves both better energy and delay performance, as stochastic routing does in these examples.

IX. CONCLUSION

A. Summary

We have presented a network routing problem which uses a probabilistic local broadcast model for wireless transmission. We then presented results showing that an index policy is optimal for this problem. We extended the model to allow for transmission control, and showed that the index nature of the optimal policy remains unchanged. Finally, we presented three distinct asynchronous distributed algorithms which compute an optimal policy, and provided results on each algorithm's convergence properties.

B. Future Research Directions

Future work in this area can take many forms. We summarize potential research on routing in ad hoc networks beyond that presented in this paper.

Time-Varying System

Quality of Service aspects of packet delivery can be included in the network model by allowing the network parameters to vary with time. A partial characterization and solution of the time-varying problem can be found in [17].

Estimation of Transmission Probabilities

The models of Section II assumed knowledge of the transmission probabilities. To actually implement the algorithms, methods to estimate these probabilities must be investigated. There is a complex interaction between the behavior of these estimation methods and how the resulting probabilities are used in the network. Numerous relevant questions arise, such as the dynamic behavior of the estimation procedure, and to what extent physical modeling of the transmission system should be incorporated into the estimation technique (as opposed to purely measurement-based approaches).

Parameter Sensitivity Analysis

Estimation of channel transmission probabilities introduces some error in the values used for routing. Other system parameters may also be only estimates, such as transmission energy cost or packet reward value. We know from Theorem 2.1 that the optimal policy for Problem (P_1) is an index policy for any value of these system parameters. But the node indices computed using system parameters in error will be different from the actual optimal indices. It is of great practical interest to determine how sensitive node indices are to errors in these system parameters. Initial results on sensitivity analysis can be found in [12].

Markov Chain Channel

It is interesting to ask what happens when the channel model is no longer a fixed iid transition probability structure, but is modeled instead by a Markov Chain. This allows for correlation in time of transition events at the same node, allowing accommodation of more general channel models. The Markov Chain channel model in general leads to a more complicated optimal policy than the Index Policy of Section II-C.

Multidestination

Problem (P_3) defines the destination to be any one node from a set of nodes, the *anycast* problem. Another important problem is the *multicast* problem, which is to send the packet to *all* the nodes in a given destination set.

General Distributed Algorithms

Algorithm 2, DDP, and Algorithm 3 are three ways to compute the optimal policy for Problem (P_1) in a distributed fashion. There are more such algorithms, and it would be useful to relate these results to a general theory of distributed convergence.

Distributed Algorithm Properties

It would be useful to characterize the convergence behavior for the algorithms presented in this paper, such as convergence rate, the counting to infinity problem, and packet looping ([15] p.252).

X. ACKNOWLEDGEMENTS

This research was supported in part by ARO Grant DAAH04-96-1-0377, AFOSR Grants F49620-96-1-0028 and F49620-98-1-0370, NSF Grants ECS-9979347 and CCR-0082784, and ONR grant N00014-03-1-0232.

REFERENCES

- [1] "1xEV-DO Evaluation Methodology (v1.4)", 3GPP2 Technical Specification Group C, C30-DOAH-20031020-0004, 2003
- [2] I. AKYILDIZ, W. SU, Y. SANKARASUBRAMANIAM, E. CAYIRCI, "A Survey on Sensor Networks", *IEEE Communications Magazine*, August 2002
- [3] D. BERTSEKAS, "Distributed Dynamic Programming", *IEEE Transactions on Automatic Control* **27**:3, p.610-616, June 1982
- [4] D. BERTSEKAS AND R. GALLAGER, *Data Networks*, Prentice-Hall, 1992
- [5] J. BROCH, D. MALTZ, D. B. JOHNSON, Y. C. HU, J. JETCHEVA, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols", *Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking*, October 25-30, 1998
- [6] M. CORSON AND A. EPHREMIDES, "A Distributed Routing Algorithm for Mobile Wireless Networks", *Wireless Networks* **1**, p. 61-81, 1995
- [7] M. FAN, ET. AL, "On the Reverse Link Performance of cdma2000 1xEV-DO Revision A System", ICC 2005
- [8] E. GAFNI AND D. BERTSEKAS "Distributed Algorithms for Generating Loop-Free Routes in Networks with Frequently Changing Topology", *IEEE Transactions on Communications* **29**:1, p.11-18, 1981
- [9] J. GARCIA-LUNA-ACEVES AND S. MURTHY "A Path-Finding Algorithm for Loop-Free Routing", *IEEE/ACM Transactions on Networking*, v 5 n 1, p.148-160, Feb 1997
- [10] M. GAREY AND D. JOHNSON *Computers and Intractability : a Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979
- [11] Z. HAAS "Guest Editorial, Wireless Ad Hoc Networks", *IEEE Journal on Selected Areas in Communications* **17**:8, p.1329-1330, August 1999
- [12] T. JAVIDI AND D. TENEKETZIS, "Sensitivity Analysis of an Optimal Routing Policy in an Ad Hoc Wireless Network", To appear in *IEEE Transactions on Automatic Control*,
- [13] D. B. JOHNSON, "Routing in Ad Hoc Networks of Mobile Hosts", *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications*, p.158-163, December 1994
- [14] G. KLIMOV, "Time Sharing Service Systems I", *Theory of Probability and its Applications*, **19**, p.532-551, 1974
- [15] J. KUROSE AND K. ROSS, *Computer Networking*, Addison-Wesley, 2000
- [16] C. LOTT, ET. AL, "Reverse Traffic Channel Design of cdma2000 1xEV-DO Revision A System", VTC June 2005
- [17] C. LOTT AND D. TENEKETZIS, "Stochastic Routing in Ad Hoc Wireless Networks", University of Michigan Control Group Report CGR 01-01, Feb. 2001
- [18] G. S. MALKIN, "RIP Version 2: Carrying Additional Information", RFC 1388 Jan, 1993
- [19] D. MALTZ, J. BROCH, J. JETCHEVA, D. JOHNSON "The Effects of On-Demand Behavior in Routing Protocols for Multihop Wireless Ad Hoc Networks", *IEEE Journal on Selected Areas in Communications* **17**:8, p.1439-1453, August 1999
- [20] P. MERLIN AND A. SEGALL "A Failsafe Distributed Routing Protocol", *IEEE Transactions on Communications* **27**:9, p.1280-1287, September 1979
- [21] J. MOY, "OSPF Version 2", RFC 1247 July, 1991
- [22] V. PARK AND M. CORSON, "A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks", *Proceedings of INFOCOM '97*, p.1405-1413, April 1997
- [23] M. PEARLMAN AND Z. HAAS "Determining the Optimal Configuration for the Zone Routing Protocol", *IEEE Journal on Selected Areas in Communications* **17**:8, p.1395-1414, August 1999
- [24] C. PERKINS AND P. BHAGWAT "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers", *Proceedings of the SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications*, p. 234-244, August 1994
- [25] C. PERKINS, E. ROYER, S. DAS, "Ad Hoc On-Demand Distance Vector (AODV) Routing", IETF MANET Working Group, March 2000.
- [26] M. PURSLEY AND H. RUSSELL "Network Protocols for Frequency-Hop Packet Radios with Decoder Side Information", *IEEE Journal on Selected Areas in Communications*, p.612-621, May 1994
- [27] M. PURSLEY, H. RUSSELL, P. STAPLES "Routing for Multimedia Traffic in Wireless Frequency-Hop Communication Networks", *IEEE Journal on Selected Areas in Communications* **17**:5, p.784-792, May 1999
- [28] V. RODOPLU AND T. MENG "Minimum Energy Mobile Wireless Networks", *IEEE Journal on Selected Areas in Communications* **17**:8, p.1333-1344, August 1999
- [29] S. ROSS, *Introduction to Stochastic Dynamic Programming*, Academic Press, 1983
- [30] W. RUDIN, *Principles of Mathematical Analysis*, McGraw-Hill, 1976
- [31] P. SAMAR, M. PEARLMAN, Z. HAAS "Independent Zone Routing: An Adaptive Hybrid Routing Framework for Ad Hoc Wireless Networks", *IEEE/ACM Transactions on Networking*, Aug. 2004
- [32] A. VITERBI, *CDMA: Principles of Spread Spectrum Communications*, Addison-Wesley 1995
- [33] J. WIESELTHIER, G. NGUYEN, A. EPHREMIDES, "Multicasting in Energy-Limited Ad-Hoc Wireless Networks" *Proceedings of the 1998 IEEE Military Communications Conference*, v.3, p.723-729, 1998
- [34] J. WIESELTHIER, G. NGUYEN, A. EPHREMIDES, "Algorithms for Bandwidth-Limited Energy-Efficient Wireless Broadcasting and Multicasting", *Proceedings of the 2000 IEEE Military Communications Conference*, Los Angeles, October 2000
- [35] H. WILF, *Generatingfunctionology*, 2nd ed, Academic Press, 1994

Appendix A - Proof of Lemma 2.6

Proof: Let $i \in \Omega$, and let $\tilde{\pi}$ be any index policy with the property that $\tilde{\pi}$ transmits when in state $\{i\}$. We have for $\tilde{\pi}$

$$\begin{aligned} V_i^{\tilde{\pi}} &= -c_i + \sum_{S \supseteq \{i\}} P^i(S|i) V_{\tilde{\pi}(S)}^{\tilde{\pi}} \\ &= -c_i + \sum_{S \supseteq \{i\}: \tilde{\pi}(S) \neq i} P^i(S|i) V_{\tilde{\pi}(S)}^{\tilde{\pi}} + \sum_{S \supseteq \{i\}: \tilde{\pi}(S) = i} P^i(S|i) V_{\tilde{\pi}(S)}^{\tilde{\pi}} \end{aligned} \quad (92)$$

Solving (92) for $V_i^{\tilde{\pi}}$ we obtain

$$V_i^{\tilde{\pi}} = \frac{-c_i + \sum_{S \supseteq \{i\}: \tilde{\pi}(S) \neq i} P^i(S|i) V_{\tilde{\pi}(S)}^{\tilde{\pi}}}{\sum_{S \supseteq \{i\}: \tilde{\pi}(S) \neq i} P^i(S|i)} \quad (93)$$

Consequently, by (41) and (93), for optimal index policy π we have

$$V_i^{\pi} = \max \left\{ \max_{\tilde{\pi}} \left\{ \frac{-c_i + \sum_{S \supseteq \{i\}: \tilde{\pi}(S) \neq i} P^i(S|i) V_{\tilde{\pi}(S)}^{\tilde{\pi}}}{\sum_{S \supseteq \{i\}: \tilde{\pi}(S) \neq i} P^i(S|i)} \right\}, R_i \right\} \quad (94)$$

Because (94) is true for any $i \in \Omega$, this proves (42). ■

Appendix B - Proof of Lemma 5.4

Proof: By Corollary 5.1, we only need to show existence of an event n where \underline{D}_i equals V_i^{π} , $\forall i \in \Omega$. By Corollary 5.1, once a node reaches V_i^{π} , it does not change thereafter. Define G to be the set of nodes at event n for which (i) $\underline{D}_i^n = V_i^{\pi}$, $i \in G$, and (ii) each $i \in G$ has successfully transmitted its optimal value V_i^{π} to all of its neighbors.

We proceed by induction. At $n = 0$ the destination nodes of highest reward have the correct V 's (i.e. R_{max}). Because transmission occurs infinitely often, there is an event at which these nodes have transmitted their values to their neighbors, and set G contains the destination nodes of reward R_{max} . This proves the induction basis step.

Now assume that at event $n - 1$ set G includes the $g > 0$ best nodes according to π , but not the $(g + 1)$ 'th best node. Let i be the $(g + 1)$ 'th best node according to π . Also assume that i recomputes its V value at event n . Event n exists because node recomputations occur infinitely often. By Lemma 5.3 (2.), the monotonicity property, and the form of the actual optimal value for i given in Lemma 2.6, the above algorithm computing the right-hand-side of (50) will give V_i^{π} . This is because neighbors not better than i according to π will have current V 's (expressed by the corresponding \underline{D} 's) less than or equal to their optimal values, which are less than i 's. Hence computation (50) in Method 1 will produce the optimal value for i . Once i successfully transmits its value after event n to all its neighbors, which occurs in finite time because transmissions occur infinitely often, node i will enter G .

By induction all nodes $i \in \Omega$ enter G in this way. Let $n_p < \infty$ be the event at which the last node enters G . This n_p satisfies the requirement of the lemma. ■

Appendix C - Proof of Lemma 5.5

Proof: By Lemma 5.3 we have that \bar{D}_i^n is monotonically non-increasing and lower bounded by $V_i^\pi, \forall i$. Hence, for each i there is a limit point of the sequence, which we label W_i . This proves 1.

To prove 2., let $\epsilon = \min_{i \neq j} \{|W_i - W_j|\}$. From 1. of Lemma 5.5 it follows that there exists an event m s.t. $\bar{D}_i^n - W_i < \epsilon, \forall n \geq m, \forall i \in \Omega$. This means at every node the node ranking (i.e. their relative priority) is fixed $\forall n \geq m$. Call $\tilde{\pi}$ the policy that uses this node ranking. As the nodes compute their values in the order of $\tilde{\pi}$, and successfully transmit them afterwards, the expected reward value $V_i^{\tilde{\pi}}$ is computed at each node i . This computation of $V_i^{\tilde{\pi}}$ at each node $i \in \Omega$ occurs in finite time because computations and communication occur infinitely often. The value $V_i^{\tilde{\pi}}$ corresponds to the expected reward at node i under $\tilde{\pi}$, a viable, possibly sub-optimal, priority. At a node i once $V_i^{\tilde{\pi}}$ is computed this value is fixed thereafter, so we must have $V_i^{\tilde{\pi}} = W_i$. We thus have

$$V_i^{\tilde{\pi}} = W_i \geq V_i^\pi, \forall i \quad (95)$$

The inequality in (95) follows from Lemma 5.3, 2. Since π is an optimal policy, $V_i^\pi \geq V_i^{\tilde{\pi}}$. Hence $\tilde{\pi}$ is also an optimal policy, which is determined in finite time. ■

Appendix D - Proof of Lemma 5.6

Proof: We use induction on i and n . First is induction on n .

Since both systems start at the same values, we have $\bar{D}_i^{v_0^i} = W_i^{u_0^i}, \forall i$. Now assume for a given n we have

$$\bar{D}_i^{v_m^i} \leq W_i^{u_m^i}, \forall m < n, \forall i \quad (96)$$

We now use induction on i . When $W_1^{u_1^n}$ is computed, all neighbor values used are of the form $W_j^{u_{n-1}^j}, j \neq 1$. When $\bar{D}_1^{v_1^n}$ is computed, all neighbor values used are of the form $\bar{D}_j^k, j \neq 1$, where $k \geq v_{n-1}^j$ (i.e. more computations beyond the $(n-1)$ 'st might have occurred in the actual (asynchronous) system). Now Lemma 5.3 and (96) imply that $\bar{D}_j^k \leq \bar{D}_j^{v_{n-1}^j} \leq W_j^{u_{n-1}^j}, \forall k \geq v_{n-1}^j$. Hence, the monotonicity property gives that $\bar{D}_1^{v_1^n} \leq W_1^{u_1^n}$. This is the first induction step on i .

Assume next that for a given i we have

$$\bar{D}_l^{v_l^n} \leq W_l^{u_l^n}, \forall l < i \quad (97)$$

When $W_i^{u_i^n}$ is computed, all neighbor values used are one of the two forms $W_j^{u_j^n}, j < i$ and $W_j^{u_{n-1}^j}, j > i$. When $\bar{D}_i^{v_i^n}$ is computed, all neighbor values used are one of the two forms $\bar{D}_j^k, j < i$, where $k \geq v_{n-1}^j$, and $\bar{D}_j^k, j > i$, where $k \geq v_{n-1}^j$ (again more computations beyond the $(n-1)$ 'st or n 'th might have occurred in the asynchronous system). We have by Lemma 5.3 and the inductive assumptions (both (96) and (97))

$$\begin{aligned} \bar{D}_j^k &\leq \bar{D}_j^{v_{n-1}^j} \leq W_j^{u_{n-1}^j}, \forall k \geq v_{n-1}^j, j < i \\ \bar{D}_j^k &\leq \bar{D}_j^{v_{n-1}^j} \leq W_j^{u_{n-1}^j}, \forall k \geq v_{n-1}^j, j > i \end{aligned}$$

So again the monotonicity property gives $\bar{D}_i^{v_i^n} \leq W_i^{u_i^n}$. This completes the induction step for i , so we have that $\bar{D}_i^{v_i^n} \leq W_i^{u_i^n}, \forall i$. And this completes the induction step for n , which completes the proof of the lemma. ■

Appendix E - Proof of Lemma 5.7

Proof: Let j be any neighbor of i . Fix all the components of $\mathbf{W}^{k,i-}$ except for that of j , where for ease of notation we assume $j < i$. The following arguments with the obvious notation change prove the result for $j > i$ as well. Because the W_l values are fixed in this proof for $l \neq i, j$, for notational convenience we write W_l for W_l^{k+1} when $l < i$, and W_l for W_l^k when $l > i$. Define

$$\mathbf{W}^{ij-} := (W_1, W_2, \dots, W_{j-1}, W_{j+1}, \dots, W_{i-1}, W_{i+1}, \dots, W_N)$$

Since the components of $\mathbf{W}^{k,i-}$ are assumed fixed except for j , we can think of f_i as being parameterized by \mathbf{W}^{ij-} , with only one argument W_j as a free variable. We then write this function as $f_i^{\mathbf{W}^{ij-}}(W_j)$. Note then that we can write (63) as

$$f_i^{\mathbf{W}^{ij-}}(W_j) = A_{\mathbf{W}^{i-}} + B_{\mathbf{W}^{i-}} \cdot W_j \quad (98)$$

where $A_{\mathbf{W}^{i-}}$ and $B_{\mathbf{W}^{i-}}$ are defined as

$$A_{\mathbf{W}^{i-}} := \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i, j} P^i(S|i) W_{\pi(S)}}{\sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i, j} P^i(S|i)} \quad (99)$$

$$B_{\mathbf{W}^{i-}} := \frac{\sum_{S \subseteq \mathcal{N}(i): \pi(S) = j} P^i(S|i)}{\sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i, j} P^i(S|i)} \quad (100)$$

Note that implicit in the definitions of $A_{\mathbf{W}^{i-}}$ and $B_{\mathbf{W}^{i-}}$ is the local policy π , which is also a function of $\mathbf{W}^{k,i-}$ (which includes W_j). But over any range of values of W_j for which π doesn't change, $A_{\mathbf{W}^{i-}}$ and $B_{\mathbf{W}^{i-}}$ are constant. Thus, over such a range $f_i^{\mathbf{W}^{ij-}}(W_j)$ is a linear function of W_j . Further, from (100) and the fact that

$$\{S \subseteq \mathcal{N}(i) : \pi(S) = j\} \subseteq \{S \subseteq \mathcal{N}(i) : \pi(S) \neq i\}, \quad (101)$$

we have that

$$0 \leq B_{\mathbf{W}^{i-}} \leq 1 \quad (102)$$

As W_j varies at values for which $W_j < f_i^{\mathbf{W}^{ij-}}(W_j)$, node j is given lower priority than i by π , and hence the value of W_j has no effect on $f_i^{\mathbf{W}^{ij-}}(W_j)$. As W_j varies at values for which $W_j \geq f_i^{\mathbf{W}^{ij-}}(W_j)$, policy π can change at only two types of points.

- 1) $W_j = f_i^{\mathbf{W}^{ij-}}(W_j)$.
- 2) $W_j = W_k$, for some $k \neq i, j$, and not type 1

A type 1 point is where node j flips priority with node i . Type 2 points are where node j flips priority with node k , but not with node i .

We first examine type 1. Let $W^* = f_i^{\mathbf{W}^{ij-}}(-\infty)$. W^* is the value when π ranks j lower than i . We then have $\forall W_j \leq W^*$ that

$$f_i^{\mathbf{W}^{ij-}}(W_j) = W^* = \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i, j} P^i(S|i) W_{\pi(S)}}{\sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i, j} P^i(S|i)} \quad (103)$$

and so the function is constant (i.e. the slope is 0).

Consider a point W_j in a linear region, as in (98) above with a fixed π . We then have $f_i^{\mathbf{W}^{ij-}}(W_j) = A_{\mathbf{W}^{i-}} + B_{\mathbf{W}^{i-}} \cdot W_j$. Let U^* be any point on this curve such that $f_i^{\mathbf{W}^{ij-}}(U^*) = U^*$, if such a point exists. We thus have $A_{\mathbf{W}^{i-}} + B_{\mathbf{W}^{i-}} \cdot U^* = U^*$, which simplifies to

$$U^* = \frac{A_{\mathbf{W}^{i-}}}{1 - B_{\mathbf{W}^{i-}}} = \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i, j} P^i(S|i) W_{\pi(S)}}{\sum_{S \subseteq \mathcal{N}(i): \pi(S) \neq i, j} P^i(S|i)} = W^* \quad (104)$$

where the final equality follows from (103). That is, a point of type 1 occurs when $W_j = W^*$, the value function for i when j is of lower priority than i in π . From (103) and (104) we conclude that such a U^* exists and is the unique point where a line of the form (98) intersects the line of slope 0 at W^* . Hence, continuity and piecewise linearity with non-decreasing slope is maintained in this case.

Now we examine a point of type 2. We only need to consider nodes k of priority higher than i under the current π (i.e. $k >_\pi i$) since j is assumed to be at a higher priority than i . Let $A'_{\mathbf{W}^{i-}}$ and $B'_{\mathbf{W}^{i-}}$ be the intercept and slope of (98) for the case where $W_j = W_j$ is slightly less than W_k , and let π' be the policy in this case. Similarly, let $A''_{\mathbf{W}^{i-}}$ and $B''_{\mathbf{W}^{i-}}$ be these values for the case where $W_j = W_j$ is slightly greater than W_k , and let π'' be this policy. We need to show that the two lines, described by (98), $A'_{\mathbf{W}^{i-}}, B'_{\mathbf{W}^{i-}}$, and (98), $A''_{\mathbf{W}^{i-}}, B''_{\mathbf{W}^{i-}}$, intersect at the point where $W_j = W_k$.

Since $j, k >_\pi i$, we have

$$\sum_{S \supset \{j\}: \pi'(S) \neq i} P^i(S|i) = \sum_{S \supset \{j\}: \pi''(S) \neq i} P^i(S|i) := G \quad (105)$$

Relation (105) is valid because π' and π'' differ only in that they switch the respective priorities of j and k . Hence, the set of states above i in priority remain identical.

We can then write (98) as

$$f_i^{\mathbf{W}^{ij-}}(W'_j) = \frac{1}{G}(-c_i + \sum_{S \subseteq \mathcal{N}(i): \pi'(S) \neq i, j, k} P^i(S|i)W_{\pi'(S)} + \sum_{S \subseteq \mathcal{N}(i): \pi'(S)=j} P^i(S|i)W'_j + \sum_{S \subseteq \mathcal{N}(i): \pi'(S)=k} P^i(S|i)W_k) \quad (106)$$

and similarly for W''_j .

From the definitions of π' and π'' , we have

$$\sum_{S \subseteq \mathcal{N}(i): \pi'(S) \neq i, j, k} P^i(S|i) = \sum_{S \subseteq \mathcal{N}(i): \pi''(S) \neq i, j, k} P^i(S|i) \quad (107)$$

Hence,

$$f_i^{\mathbf{W}^{ij-}}(W''_j) - f_i^{\mathbf{W}^{ij-}}(W'_j) = \frac{1}{G}(\sum_{S \subseteq \mathcal{N}(i): \pi'(S)=j} P^i(S|i)W'_j + \sum_{S \subseteq \mathcal{N}(i): \pi'(S)=k} P^i(S|i)W_k - \sum_{S \subseteq \mathcal{N}(i): \pi''(S)=j} P^i(S|i)W''_j - \sum_{S \subseteq \mathcal{N}(i): \pi''(S)=k} P^i(S|i)W_k) \quad (108)$$

When $W'_j = W''_j = W_k$, (108) becomes

$$f_i^{\mathbf{W}^{ij-}}(W''_j) - f_i^{\mathbf{W}^{ij-}}(W'_j) = \frac{W_k}{G}(\sum_{S \subseteq \mathcal{N}(i): \pi'(S)=j} P^i(S|i) + \sum_{S \subseteq \mathcal{N}(i): \pi'(S)=k} P^i(S|i) - \sum_{S \subseteq \mathcal{N}(i): \pi''(S)=j} P^i(S|i) - \sum_{S \subseteq \mathcal{N}(i): \pi''(S)=k} P^i(S|i)) = 0 \quad (109)$$

where the terms in parenthesis sum to 0 because π' and π'' are identical except for the switching of the adjacent priorities of j and k , so that the sum over all states where either of these two nodes are the highest priority is the same for both π' and π'' .

Because the linear functions described by (98), $A_{\mathbf{W}^{i-}}$, $B_{\mathbf{W}^{i-}}$, and (98), $A_{\mathbf{W}^{i-}}$, $B_{\mathbf{W}^{i-}}$, intersect at $W'_j = W''_j = W_k$, $f_i^{\mathbf{W}^{ij-}}(W_j)$ is piecewise linear over the region of W_j where π' and π'' are used. Furthermore, $B_{\mathbf{W}^{i-}} \geq B_{\mathbf{W}^{i-}}$ because the denominators are the same in both cases, and $B_{\mathbf{W}^{i-}}$ has a larger numerator than $B_{\mathbf{W}^{i-}}$, because $j >_{\pi''} k$ and $j <_{\pi'} k$. Hence the slope is monotonically non-decreasing in this region. ■

Appendix F - Proof of Lemma 5.9

Proof: By Corollary 5.1, \mathbf{V}^π is a fixed point of T .

Now assume there exists $\mathbf{W} \neq \mathbf{V}^\pi$ s.t. $\mathbf{W} = T(\mathbf{W})$. Let

$$i \in \argmax_{1 \leq j \leq N} \{W_j : W_j \neq V_j^\pi\} \quad (110)$$

This i may not be unique (there may be a tie), but if so, the following still holds for any one of the maxima. By Lemma 5.3 and Lemma 5.4 we have $W_i \geq V_i^\pi$. This fact and (110) give

$$W_i > V_i^\pi \quad (111)$$

Let $S := \{j \in \Omega : W_j > V_j^\pi\}$. Note that S cannot be empty, because S empty means $W_i = R_i \leq V_i^\pi$, contradicting (111). So S is not empty, and by (110) we have

$$W_j = V_j^\pi, \forall j \in S \quad (112)$$

This fact, the definition of $T(\cdot)$, and (63) imply that $T(W)$ uses only V_j^π values in the computation of $W_i = T(\mathbf{W})|_i$. Let π^i be the local policy for i determined by Method 1, and note that π^i ranks only the nodes of S above i . Then we have by Lemma 2.6 that

$$W_i = T(\mathbf{W})|_i = V_i^{\pi^i} \leq V_i^\pi \quad (113)$$

The final inequality of (113) follows because of (112) and the fact that π has an optimal local ordering.

Relation (113) contradicts (111). Hence, no such i can exist, and this completes the proof of the lemma. ■

Appendix G - Proof of Lemma 7.1

Proof: There is no requirement that the node rankings of $\hat{\pi}$ and $\hat{\pi}$ reflect the actual current ranking of neighbor values $\{V_{i,j}\}$.

The update computation of (84) at node i for local policy $\hat{\pi}$ is

$$V_i^{\hat{\pi}} = \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \neq i} P^i(S|i)V_{i, \hat{\pi}(S)}}{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \neq i} P^i(S|i)} \quad (114)$$

where to ensure (114) is well defined we assume that

$$\sum_{\{i\} \subset S \subseteq \mathcal{N}(i): \hat{\pi}(S) \neq i} P^i(S|i) \neq 0 \quad (115)$$

A similar computation holds for local policy $\hat{\pi}$.

$$V_i^{\hat{\pi}} = \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \neq i} P^i(S|i)V_{i, \hat{\pi}(S)}}{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \neq i} P^i(S|i)} \quad (116)$$

Equation (116) is well defined because (115) also ensures its denominator is non-zero.

Write (116) as

$$V_i^{\hat{\pi}} = \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i)V_{i, \hat{\pi}(S)}}{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i) + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in B} P^i(S|i)} + \frac{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in B} P^i(S|i)V_{i, \hat{\pi}(S)}}{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i) + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in B} P^i(S|i)}$$

Using the definition of Y_1 and (117) we have

$$V_i^{\hat{\pi}} \leq \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i)V_{i, \hat{\pi}(S)}}{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i) + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in B} P^i(S|i)} + \frac{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in B} P^i(S|i)Y_1}{\sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i) + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in B} P^i(S|i)} \quad (117)$$

Note that since both $\hat{\pi}$ and $\hat{\pi}$ make A the highest priority nodes with the same ranking, we have

$$\{S \subseteq \mathcal{N}(i) : \hat{\pi}(S) \in A\} = \{S \subseteq \mathcal{N}(i) : \hat{\pi}(S) \in A\} \quad (118)$$

$$\hat{\pi}(S) = \hat{\pi}(S), S \subseteq \mathcal{N}(i) : \hat{\pi}(S) \in A(119)$$

For simplicity define

$$b_1 := \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i) = \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i) \quad (120)$$

$$b_2 := \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in B} P^i(S|i) \quad (121)$$

where the equality in (120) follows from (118). Because of (115), $b_1 > 0$; furthermore

$$0 < b_1 + b_2 \leq 1 \quad \text{and} \quad b_2 \geq 0 \quad (122)$$

Because of (119) and (120), (117) gives

$$V_i^{\hat{\pi}} \leq \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i)V_{i, \hat{\pi}(S)} + b_2 Y_1}{b_1 + b_2} \quad (123)$$

Write (123) as

$$V_i^{\hat{\pi}} \leq \left(\frac{b_1}{b_1 + b_2} \right)^{-c_i + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i)V_{i, \hat{\pi}(S)}}_{b_1} + \left(\frac{b_2}{b_1 + b_2} \right) Y_1$$

$$= \left(\frac{b_1}{b_1 + b_2} \right)^{-c_i + \sum_{S \subseteq \mathcal{N}(i): \hat{\pi}(S) \in A} P^i(S|i)V_{i, \hat{\pi}(S)}}_{b_1} + \left(\frac{b_2}{b_1 + b_2} \right) Y_1 \quad (124)$$

$$= \left(\frac{b_1}{b_1 + b_2} \right) V_i^{\hat{\pi}} + \left(\frac{b_2}{b_1 + b_2} \right) Y_1 \quad (125)$$

Equation (124) follows from (119), and (125) follows from (114).

Define

$$a := \frac{b_2}{b_1 + b_2} \quad (126)$$

Since $b_1 > 0$ and $b_2 \geq 0$, we have $0 \leq a < 1$. Inequality (125) is then

$$V_i^\pi \leq (1 - a) V_i^\pi + a Y_1 \quad (127)$$

A similar argument using Y_2 gives

$$V_i^\pi \geq (1 - a) V_i^\pi + a Y_2 \quad (128)$$

Appendix H - Proof of Lemma 7.2

Proof: We have

$$V_i^n = \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i) V_{i, \tilde{\pi}(S)}^n}{\sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i)} \quad (129)$$

$$\leq \frac{-c_i + \sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i) V_{max}}{\sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i)} \quad (130)$$

$$= V_{max} - \frac{c_i}{\sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i)} \quad (131)$$

$$\leq V_{max} - c_i \quad (132)$$

where (132) follows because $\sum_{S \subseteq \mathcal{N}(i): \tilde{\pi}(S) \neq i} P^i(S|i) < 1$ and $c_i > 0$. Inequality (132) completes the the proof of the lemma. ■

Appendix I - Proof of Lemma 7.3

Proof: Because π_j is optimal, when

$$V_{j,k} = V_k^\pi, \quad \forall k \in \mathcal{N}(j) \quad (133)$$

by Lemma 2.6 (42) we have that

$$V_j^\pi = \max \left\{ \frac{-c_j + \sum_{S \subseteq \mathcal{N}(j): \pi_j(S) \neq j} P^j(S|j) V_{j, \pi_j(S)}^\pi}{\sum_{S \subseteq \mathcal{N}(j): \pi_j(S) \neq j} P^j(S|j)}, R_j \right\} \quad (134)$$

Hence, under (133) the update (84) using π_j computes the value function V_j^π . By [30] Section 4.11, every rational function $f: \mathbb{R}^N \rightarrow \mathbb{R}$ is continuous. Because: 1) Update (84) is a composition of a rational function of the values $V_{j,k}, k \in \mathcal{N}(j)$, with the function $\max\{\cdot, \cdot\}$; 2) the function $\max\{\cdot, \cdot\}$ is continuous in its arguments, 3) the composition of continuous functions results in a continuous function, and 4) only node values $V_{j,k}, k \in B$ affect the (84) update using π_j , we conclude that update (84) using π_j is a continuous function mapping $\mathbb{R}^l \rightarrow \mathbb{R}$, where $l := |B|$. Since (84) is monotonic in $V_{j,k}$, the result (89) follows.

To obtain (90), we note that under (133) for any local policy $\tilde{\pi}_j$ we have

$$V_j^{\tilde{\pi}_j} \leq V_j^\pi \quad (135)$$

By the same argument given above, for fixed $\tilde{\pi}_j$ the update (84) using $\tilde{\pi}_j$ is a continuous function mapping $\mathbb{R}^l \rightarrow \mathbb{R}$, where $l := |\bar{B}|$, and where only node values $V_{j,k}, k \in \bar{B}$ affect the update. Because (84) is monotonic in $V_{j,k}$, the result (90) follows. ■

Appendix J - Proof of Lemma 7.4

Proof: We proceed by induction on the number of nodes. Assume we have a set of nodes $D \subset \Omega, D \neq \emptyset$ with the following properties:

$$\lim_{n \rightarrow \infty} V_i^n = V_i^\pi, \quad \forall i \in D \quad (136)$$

$$V_i^\pi > V_j^\pi, \quad \forall i \in D, j \in \bar{D} \quad (137)$$

where $\bar{D} := \Omega - D$.

Furthermore, because of (136) there exists an event m_1 and $\delta > 0$ such that

$$V_i^n \leq V_i^\pi + \delta, \quad \forall n > m_1, \forall i \in D \quad (138)$$

For the induction basis we let $D = \{i \in \Omega : R_i = R_{max}\}$, and argue as follows. Equation (136) is satisfied because $V_i^n = R_{max}, \forall n, i \in D$; inequality (137) is true because $c_j > 0, \forall j \in \Omega$ and hence $V_j^\pi < R_{max}, j \in \bar{D}$.

We proceed with the induction step. Define

$$E := \operatorname{argmax}_{i \in \bar{D}} V_i^\pi \quad (139)$$

Note that E is a non-empty set of nodes. Let node $i \in E$.

Claim 1:

$$\limsup_{n \rightarrow \infty} V_i^n \leq V_i^\pi \quad (140)$$

■ *Proof of Claim 1:* See Appendix K. ◀

Claim 2:

$$\liminf_{n \rightarrow \infty} V_i^n \geq V_i^\pi \quad (141)$$

Proof of Claim 2: See Appendix L. ◀

From Claim 1 and Claim 2 it follows that

$$\lim_{n \rightarrow \infty} V_i^n = V_i^\pi \quad (142)$$

The arguments leading to (142) can be made for any element $i \in E$. Hence we have shown that

$$\lim_{n \rightarrow \infty} V_i^n = V_i^\pi, \quad \forall i \in E \quad (143)$$

We now complete the induction step. Let $D' := D \cup E$ and $\bar{D}' := \Omega - D'$. We show that (136) and (137) are satisfied for D' and \bar{D}' .

Relation (143) and the induction assumption together mean that (136) is true for all nodes of D' . The definition of E in (139) with the induction assumption directly implies that (137) is true for D' .

This completes the induction argument, and the lemma is proved. ■

Appendix K - Proof of Claim 1

Proof: To prove (140) it suffices to show that

$$\limsup_{n \rightarrow \infty} Z_n \leq V_i^\pi \quad (144)$$

where

$$Z_n := \max \left\{ \max_{j \in \bar{D}} V_j^n, \max_{j,k \in \bar{D}} V_{j,k}^n \right\} \quad (145)$$

To prove (144) we show that for any $\epsilon > 0$ there is an event m such that

$$Z_n \leq V_i^\pi + \epsilon, \quad \forall n \geq m \quad (146)$$

Consider a node update at event n for any $j \in \bar{D}$. Define

$$A_j^n := \{k \in \mathcal{N}(j) \cap D : V_{j,k}^{n-1} > V_j^{n-1}\}, \quad j \in \bar{D} \quad (147)$$

A_j^n is the set of neighbor nodes in D which are ranked at higher priority than j in the update at n . A_j^n may or may not be empty. Define

$$F_n := \{j \in \bar{D} : A_j^n \neq \emptyset\} \quad (148)$$

$$\bar{F}_n := \{j \in \bar{D} : A_j^n = \emptyset\} \quad (149)$$

We first establish the following fact (recall that m_1 is defined in (138)).

Fact 1: When $Z_{n-1} \leq V_i^\pi + \epsilon_1, 0 < \epsilon_1 < 1$, then at event $n > m_1$,

$$Z_n \leq Z_{n-1} \quad \text{if event } n \text{ is a node transmission} \quad (150)$$

$$V_j^n \leq Z_{n-1} \quad \text{if event } n \text{ is an update at node } j, j \in \bar{F}_n \quad (151)$$

$$V_j^n \leq V_i^\pi + \epsilon_1 \quad \text{if event } n \text{ is an update at node } j, j \in F_n \quad (152)$$

Proof of Fact 1: If event n is a node transmission, then (150) follows from the definition of Z_n .

If event n is an update at node $j, j \in \bar{F}_n$, we consider three cases:

1) If the update gives

$$V_j^n = R_i \quad (153)$$

then

$$V_j^n = R_i \leq V_j^{n-1} \leq Z_{n-1} \quad (154)$$

2) If the update is a reset for $j \in \bar{F}_n$, then

$$V_j^{n-1} \geq V_{j,k}^{n-1}, \forall k \in \mathcal{N}(j) \quad (155)$$

Upon resetting, by the specification of Algorithm 3 node j is ranked lower than at least one neighbor, say $k \in \mathcal{N}(j)$. Then by Lemma 7.2 we obtain

$$V_j^n \leq V_{j,k}^{n-1} - c_j \leq V_j^{n-1} - c_j \leq Z_{n-1} - c_j \leq Z_{n-1} \quad (156)$$

3) If the update is a rank update, then from the definitions of F_n , Z_n , and Lemma 7.2 it follows that

$$V_j^n \leq \max_{k \in \bar{D} \cap \mathcal{N}(j)} V_{j,k}^{n-1} - c_j \quad (157)$$

$$\leq Z_{n-1} - c_j \leq Z_n \quad (158)$$

From (154), (156) and (158) it follows that (151) is true.

If event n is an update at node j , $j \in F_n$, then $A_j^n \neq \emptyset$ and the update is not a node reset. Let π_j^n be the index ranking local to j which ranks the nodes of A_j^n in the order of their value functions V_k^n , $k \in A_j^n$, and ranks j next. Since $A_j^n \subseteq D$ because of (138), we get

$$V_{j,k}^n \leq V_k^n + \delta, \forall n > m_1, \forall k \in A_j^n, \delta > 0 \quad (159)$$

Hence by Lemma 7.3 (inequality (90)) we get

$$V_j^{\pi_j^n} \leq V_j^n + \epsilon_1 \leq V_i^n + \epsilon_1 \quad (160)$$

We use (160) together with Lemma 7.1 to prove (152). We let $A = A_j^n$, $\tilde{\pi} = \pi_j^n$, $B \subseteq \mathcal{N}(j) \cap \bar{D}$ in Lemma 7.1. Then Lemma 7.1 and (160) give

$$V_j^n \leq (1-a)V_j^{\pi_j^n} + aY_1 \quad (161)$$

$$\leq (1-a)V_j^{\pi_j^n} + aZ_{n-1} \quad (162)$$

$$\leq (1-a)(V_i^n + \epsilon_1) + aZ_{n-1} \quad (163)$$

$$\leq V_i^n + \epsilon_1 \quad (164)$$

where

$$Y_1 := \max_{k \in B} V_{j,k}^{n-1} \quad (165)$$

and a is defined by relations similar to (120), (121), and (126).

The inequality in (161) follows from Lemma 7.1. The inequality in (162) follows from the definitions of Y_1 and Z_{n-1} . The inequality in (163) follows from (160). The inequality in (164) follows from the assumption $Z_{n-1} \leq V_i^n + \epsilon_1$ made in Fact 1.

Inequalities (161)-(164) prove (152). The proof of Fact 1 is now complete. \triangleleft

We proceed with the proof of Claim 1 by establishing the following fact.

Fact 2: When $Z_{n-1} > V_i^n + \epsilon_1$, $0 < \epsilon_1 < \epsilon$, then at event $n > m_1$,

$$Z_n \leq Z_{n-1} \quad \text{if event } n \text{ is a node transmission} \quad (166)$$

$$V_j^n \leq \max\{Z_{n-1} - c_j, R_j\} \quad \text{if event } n \text{ is an update at node } j, j \in \bar{F}_n \quad (167)$$

$$V_j^n \leq \max\{(1-a_{\max})(V_i^n + \epsilon_1) + a_{\max}Z_{n-1}, R_j\} \quad \text{if event } n \text{ is an update at node } j, j \in F_n \quad (168)$$

Proof of Fact 2: If event n is a node transmission, then (166) follows from the definition of Z_n

If event n is an update at node j , $j \in \bar{F}_n$, we consider three cases:

1) The update gives

$$V_j^n = R_i \quad (169)$$

2) The update is a reset for $j \in \bar{F}_n$. Then by the same arguments as those leading to (156) we obtain

$$V_j^n \leq Z_{n-1} - c_j \quad (170)$$

3) The update is a rank update for node $j \in \bar{F}_n$. Then by arguments that are the same as those leading to (158) we obtain

$$V_j^n \leq Z_{n-1} - c_j \quad (171)$$

From (169)-(171) it follows that

$$V_j^n \leq \max\{Z_{n-1} - c_j, R_j\} \quad (172)$$

which is precisely (167).

If event n is an update at node j , $j \in F_n$, then if the optimal action is not to retire, by arguments that are the same as those leading to (163) we obtain

$$V_j^n \leq (1-a)(V_i^n + \epsilon_1) + aZ_{n-1} \quad (173)$$

where a is again defined by relations similar to (120), (121), and (126).

Furthermore,

$$a_{\max} := \max_{i \in \Omega, \tilde{\pi}, \hat{\pi}} a(i, \tilde{\pi}, \hat{\pi}) \geq a \quad (174)$$

and

$$Z_{n-1} > V_i^n + \epsilon_1 \quad (175)$$

by assumption.

Because of (174) and (175), (173) gives

$$\begin{aligned} V_j^n &\leq (1-a)(V_i^n + \epsilon_1) + aZ_{n-1} \\ &\leq (1-a_{\max})(V_i^n + \epsilon_1) + a_{\max}Z_{n-1} \end{aligned} \quad (176)$$

If the optimal action is to retire, then

$$V_j^n = R_i \quad (177)$$

Combining (176) and (177), we obtain

$$V_j^n \leq \max\{(1-a_{\max})(V_i^n + \epsilon_1) + a_{\max}Z_{n-1}, R_j\} \quad (178)$$

which is precisely (168).

This completes the proof of Fact 2. \triangleleft

We use Fact 1 and Fact 2 to complete the proof of Claim 1. We consider two cases.

Case 1 There is an event $l > m_1$ where $Z_l \leq V_i^n + \epsilon_1$.

Under Case 1, relations (150), (151), and (152) together imply that

$$Z_n \leq V_i^n + \epsilon_1 \quad \forall n \geq l \quad (179)$$

Because inequality (146) follows from (179), we have shown that Claim 1 is true under Case 1.

Case 2 $Z_n > V_i^n + \epsilon$, $\forall n > m_1$

We claim that under Case 2 the following inequality (180) is true for any event, update or transmission, of any node $j \in \bar{D}$ at any $n > m_1$.

$$V_j^n \leq Z_{n-1}, \quad n > m_1 \quad (180)$$

We prove (180) as follows.

1) Suppose there is an update at $j \in \bar{F}_n$. Then

a) If $\max\{Z_{n-1} - c_j, R_j\} = R_j$, then from (167) we have $V_j^n \leq R_j$. Since $V_j^n \geq R_j$ by its definition, it follows that

$$V_j^n = R_j \quad (181)$$

And since $R_j \leq Z_{n-1}$, (180) follows.

- b) If $\max\{Z_{n-1} - c_j, R_j\} = Z_{n-1} - c_j$, then from (167) we have

$$V_j^n \leq Z_{n-1} - c_j < Z_{n-1} \quad (182)$$

2) Suppose there is an update at $j \in F_n$. Then

- a) If $\max\{(1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{n-1}, R_j\} = R_j$, then from (168) we have $V_j^n \leq R_j$. Since $V_j^n \geq R_j$ by its definition, it follows that

$$V_j^n = R_j \quad (183)$$

And since $R_j \leq Z_{n-1}$, (180) follows.

- b) If $\max\{(1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{n-1}, R_j\} = (1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{n-1}$, then from (168) we have

$$V_j^n \leq (1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{n-1} < (1 - a_{max})Z_{n-1} + a_{max}Z_{n-1} = Z_{n-1} \quad (184)$$

The strict inequality in (184) follows because $Z_n > V_i^\pi + \epsilon_1, \forall n > m_1$.

- 3) Suppose there is a transmission from $j \in \bar{D}$. Then (180) follows from (166). This completes the proof of (180).

From (180) it follows that

$$Z_n \leq Z_{n-1}, \quad n > m_1 \quad (185)$$

From (185) we also have

$$Z_{n-1} \leq Z_{m_1}, \quad n > m_1 \quad (186)$$

We complete the proof of Claim 1 under Case 2 by establishing the following two facts.

Fact 3: For all $n > m_1$ and $j \in \bar{D}$ we have

$$V_j^n \leq \max\{(1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{m_1}, Z_{m_1} - c_j\} \quad (187)$$

Proof of Fact 3: Using (186) in (168), we have for an update at $j \in F_n$ that

$$V_j^n \leq \max\{(1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{m_1}, R_j\}, \quad n > m_1 \quad (188)$$

Because $Z_{m_1} > V_i^\pi + \epsilon_1 \geq R_j, \forall j \in \bar{D}$, we can write (188) as

$$V_j^n \leq (1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{m_1}, \quad n > m_1 \quad (189)$$

Using (186) in (167), we obtain for $n > m_1$

$$V_j^n \leq \max\{Z_{m_1} - c_j, R_j\} \quad (190)$$

Using (189), (190) and the fact that $R_j \leq V_j^n \leq (1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{m_1}$, we have for $j \in \bar{D}$ and $n > m_1$ that

$$V_j^n \leq \max\{(1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{m_1}, Z_{m_1} - c_j\} \quad (191)$$

◁

Fact 4: There is an event $m_2 > m_1$ where

$$Z_{m_2} \leq (1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{m_1} \quad (192)$$

Proof of Fact 4: We prove (192) by contradiction. Assume (192) is not true, which means that

$$Z_n > (1 - a_{max})(V_i^\pi + \epsilon_1) + a_{max}Z_{m_1}, \quad \forall n > m_1 \quad (193)$$

By the definition of Z_n , (187), (193) and the fact that node transmissions and updates occur infinitely often (so that for all $j, k \in \bar{D}$, V_j^n satisfying (187) becomes $V_{k,j}^{n'}$ at some $n' > n$) there exists $l > m_1$ such that

$$Z_l \leq Z_{m_1} - \min_{j \in \bar{D}} c_j \quad (194)$$

By repeating the argument leading to (194), we conclude that for large enough n , Z_n can become arbitrarily small, and this contradicts (193). ◁

Repeating the argument leading to (192), and noting that $a_{max} < 1$, we conclude that there is an event m_p such that

$$Z_n \leq V_i^\pi + \epsilon_1 + \epsilon_2, \quad \forall n \geq m_p \quad (195)$$

Relation (195) proves (146) and hence Claim 1 under Case 2.

We have proved Claim 1 under both Case 1 and Case 2, and this completes the proof of Claim 1. ■

Appendix L - Proof of Claim 2

Proof: To prove (141) we show that for any $\epsilon > 0$ there is an event m such that

$$V_i^n \geq V_i^\pi - \epsilon, \quad \forall n \geq m \quad (i \in E) \quad (196)$$

If $V_i^\pi = R_i$, then (196) holds. Henceforth, assume $V_i^\pi > R_i$.

Define $A_i := \mathcal{N}(i) \cap D$. A_i is not empty, reasoned as follows. Because $V_i^\pi > R_i$, a rank update with correct neighbor values and optimal local policy gives V_i^π . Lemma 7.2 and the fact that $c_i > 0$ ensure that $V_i^\pi < V_j^\pi$ for some $j \in \mathcal{N}(i)$. Because of (137) and (139), $j \in D$. Hence, $j \in A_i$.

Let π_i be the index policy local to i which ranks the nodes of A_i in the order of their value functions $V_j^\pi, j \in A_i$, and then ranks i next. Note that due to (139), π_i is the optimal local policy for node i . We write $V_i^n|_{\pi_i}$ to indicate the update value computed at node i under local policy π_i .

Let $\epsilon > 0$, and choose any ϵ_1 and ϵ_2 such that $\epsilon_1 > 0, \epsilon_2 > 0$ and $\epsilon = \epsilon_1 + \epsilon_2$.

We establish the following fact.

Fact 5: There is an event $m_2 > m_1$ such that for $n \geq m_2$ we have

$$V_i^n \geq (1 - a_n)(V_i^\pi - \epsilon_1) + a_n V_i^{n-1} \quad (197)$$

where $0 \leq a_n \leq a_{max}$, $n \geq m_2$, m_1 is defined in (138), and a_{max} is given by (91).

Proof of Fact 5: By Lemma 7.3 (89) there is a $\delta > 0$ such that when $V_{i,j}^{n-1} \geq V_j^\pi - \delta, \forall j \in A_i$, we have

$$V_i^n|_{\pi_i} \geq V_i^\pi - \epsilon_1 \quad (198)$$

Let $m' > m_1$ be an event such that

$$V_{i,k}^n \geq V_k^\pi - \delta, \quad \forall n \geq m', k \in D \quad (199)$$

Event m' exists because of (136) and because transmissions occur infinitely often. Hence (198) holds for $n \geq m'$.

By Claim 1, (136) and (137), there is an event $m'' > m'$ such that $\forall j \in \bar{D}$ we have

$$V_j^n < V_k^n, \quad \forall n \geq m'', \forall k \in D \quad (200)$$

Because of (200) and the fact that transmissions occur infinitely often, there is an event $m_2 > m''$ such that $\forall j \in \bar{D}$ we have

$$V_j^n < V_{j,k}^n, \quad \forall n \geq m_2, \forall k \in D \quad (201)$$

Consider a node update at event $n > m_2$ at node i . To apply Lemma 7.1, we note that due to (201) an update at i will rank the nodes of A_i highest, then possibly rank other neighbors above i . We identify $A = A_i$ and let $B = \{j \in \bar{D} : V_{i,j}^{n-1} > V_i^{n-1}\}$. Define $Y_2 := \min_{j \in B} V_{i,j}^{n-1}$. By Lemma 7.1, at each update $n > m_2$ there exists $0 \leq a_n < a_{max}$ such that

$$V_i^n \geq (1 - a_n)V_i^n|_{\pi_i} + a_n Y_2 \quad (202)$$

$$\geq (1 - a_n)V_i^n|_{\pi_i} + a_n V_i^{n-1} \quad (203)$$

$$\geq (1 - a_n)(V_i^\pi - \epsilon_1) + a_n V_i^{n-1} \quad (204)$$

Relation (203) follows from and the definitions of B and Y_2 . Relation (204) follows from (198). ◁

We use Fact 5 to prove Claim 2 by considering two cases.

Case 1 There is an event $l > m_2$ where $V_i^l \geq V_i^\pi - \epsilon_1$.

When at some update event $n > m_2$ we have $V_i^{n-1} \geq V_i^\pi - \epsilon_1$, from (204) we obtain

$$V_i^n \geq V_i^\pi - \epsilon_1 \quad (205)$$

Relation (205) implies that if the Case 1 condition is met at $l > m_2$, it follows by (204) that

$$V_i^n \geq V_i^\pi - \epsilon_1, \quad \forall n \geq l \quad (206)$$

Relation (206) proves (196). Hence Claim 2 is proved for Case 1.

Case 2 $V_i^n < V_i^\pi - \epsilon_1, \forall n > m_1$

Under Case 2, we use the facts that $V_i^{n-1} < V_i^\pi - \epsilon_1, \forall n > m_2 + 1, 0 \leq a_n \leq a_{max}, \forall n > m_2$ and (204) to obtain

$$V_i^n \geq (1 - a_{max})(V_i^\pi - \epsilon_1) + a_{max}V_i^{n-1}, \quad \forall n > m_2 + 1 \quad (207)$$

Let x_i^n be defined as follows.

$$x_i^{m_2+1} = V_i^{m_2+1} \quad (208)$$

$$x_i^n = (1 - a_{max})(V_i^\pi - \epsilon_1) + a_{max}x_i^{n-1}, \quad \forall n > m_2 + 1 \quad (209)$$

Using (207), (208), and (209), and the fact that $a_{max} \geq 0$ with an inductive argument, we have that

$$V_i^n \geq x_i^n, \quad \forall n \geq m_2 + 1 \quad (210)$$

From (209) and the fact that $0 \leq a_{max} < 1$, we have that

$$\lim_{n \rightarrow \infty} x_i^n = V_i^\pi - \epsilon_1 \quad (211)$$

From (210) and (211), we conclude that

$$\lim_{n \rightarrow \infty} V_i^n \geq V_i^\pi - \epsilon_1 \quad (212)$$

Relation (212) implies that there exists an event $m_3 > m_2 + 1$ such that

$$V_i^n \geq V_i^\pi - \epsilon_1 - \epsilon_2, \quad \forall n \geq m_3 \quad (213)$$

Relation (213) proves (196). Hence Claim 2 is proved for Case 2. ■

PLACE
PHOTO
HERE

Christopher Lott Qualcomm, Inc. Formerly University of Michigan, Department of EECS

PLACE
PHOTO
HERE

Demosthenis Teneketzis Professor, University of Michigan, Department of EECS