Routing with probabilistic delay guarantees in wireless ad-hoc networks

Matthew Brand,[1] Petar Maymounkov,[2] and Andreas F. Molisch,[1] *Fellow, IEEE*
[1]Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA
email: {brand, molisch}@merl.com
[2]Massachusetts Institute of Technology (MIT), Cambridge, MA, USA
email: petar@csail.mit.edu.

***Abstract*— In many wireless ad-hoc networks it is important to find a route that delivers a message to the destination within a certain deadline (delay constraint). We propose to identify such routes based on average channel state information (CSI) only, since this information can be distributed more easily over the network. Such cases allow probabilistic QoS guarantees i.e., we maximize and report the probability of on-time delivery. We develop a convolution-free lower bound on probability of on-time arrival, and a scheme to rapidly identify a path that maximizes this bound. This analysis is motivated by a class of infinite variance subexponential distributions whose properties preclude the use of deviation bounds and convolutional schemes. The bound then forms the basis of an algorithm that finds routes that give probabilistic delay guarantees. Simulations demonstrate that the algorithm performs better than shortest-path algorithm based on average CSI.**

I. Introduction

Wireless ad-hoc networks have in recent years emerged as the most promising way to achieve ubiquitous, reliable connectivity. All nodes have similar functionality, and information is forwarded from the source to the destination via a number of other (relaying) nodes. Ad-hoc networks provide advantages to (i) cost, since no fixed infrastructure is required, (ii) flexibility and ease of deployability, and (iii) reliability, since the elimination of a single node does not lead to a failure of the whole network. For this reason ad-hoc networks are popular for applications ranging from communications for emergency responders, collection of environmental data, factory automation to security and military applications.

In many applications ad-hoc networks have to provide a guarantee for quality-of-service (QoS); in this paper we particularly consider the transmission delay from source to destination. For example, a message that a piece of machinery is overheating has to be delivered to the control center before the machine destroys itself. QoS in wireless ad-hoc networks is influenced by a wide variety of factors, among them (i) call admission, (ii) arrival statistics of packets from higher layers, (iii) scheduling and multiple-access mechanisms, (iv) properties of the physical-layer transmission, and (v) routing. For a survey of these issues and methods to deal with them, see, e.g., [1], [2]. Due to all those statistical variations, in particular of the physical layer, it is not possible to give a perfect guarantee that a packet will arrive at the receiver within a certain time; it is only possible to guarantee that in a certain percentage of all channel realizations (e.g., 99%), the packets will arrive in time. We will henceforth refer to such a statement for the probability of on-time arrival as a "probabilistic guarantee".[1] We note that while stochastic variations of the delay due to random packet arrival of the source have been treated extensively in ther literature (see [4] and references therein), random variations of the transmission time due to randomly varying channels has drawn very little attention.

A topic of particular importance for QoS in ad-hoc networks is routing, i.e., determining the nodes over which the information is forwarded from source to destination [5]. Routing algorithms can be roughly categorized as follows: (i) flooding and gossiping, where the information is sent out from the source, and either all, or randomly chosen nodes forward the information. This approach does not require any knowlege of channel state information by the nodes, but is energy inefficient. (ii) geometry-based routing algorithms: an optimum route is identified (in a central or distributed way) based on the knowledge of the location of the nodes. However, since short distance between two nodes does not necessarily mean good propagation conditions, such algorithms can lead to suboptimum routes, (iii) route determination based on instantaneous channel state information (CSI), also called "stateful approach": in this category of algorithms, the optimum route is determined from a global or distributed knowledge of the instantaneous CSI of all the links [6], [7], [8]. This category also subsumes methods that send out route discovery packets and store the results in routing tables. A route that fulfills the QoS constraints is determined and kept until it breaks (i.e., the QoS constraint is violated); then the route is either repaired locally [9], [10] or a new route is determined.

For many applications, route discovery based on instantaneous CSI is not feasible. Since wireless channel states can be constantly changing, a frequent update of the CSI throughout the network would lead to unacceptable overhead (typical coherence times of wireless propagation channels, i.e, the required update interval, is on the order of a few milliseconds [3]). Especially in large networks the overhead traffic communicating the routing information for all possible links would decrease spectral efficiency and battery lifetime. On the other hand, on-demand route discovery is not feasible because the route dis-

This work was done while all the authors were at MERL.

[1]This notion is similar in spirit to "outage probability" of cellular networks, which defines the probability that a mobile station does not receive sufficient signal power to communicate with a base station [3].

covery process often takes longer than the admissible delay of the information. For this reason, we provide in this paper a method to perform routing based on the *average* CSI. Average CSI changes only very slowly, so that it can be communicated through a network without undue overhead.[2] In particular, our contributions are as follows

- We show that routing based on the statistics of the channel state can provide probabilitistic quality-of-service guarantuees, in particular, a guarantee that packets are delivered to the destination within a deadline t in a fraction p of all channel realizations.
- We introduce an extremely efficient routing algorithm that finds the path that provides high QoS. In contrast to the mostly heuristic routing algorithms in the literature, we provide an analytical proof that the algorithm maximizes a lower bound for the probability for on-time delivery.

The remainder of the paper is organized as follows: Section II outlines the system model, in particular the assumptions about the network topology, transmission scheme over one link, and quality-of-service requirements. Section III is the core of the paper, providing the algorithm, as well as the analytical proof that it actually provides the quality of service. Section IV demonstrates the algorithm by showing some simulation results. A summary and conclusions wrap up the paper.

II. SYSTEM MODEL

We consider a wireless network with K randomly placed nodes; in Sec. III, the network will be described as a graph with K nodes and n edges, i.,e., connections between the nodes. Our goal is the transmission of a message from one source to one destination (unicast), so that the delay is no larger than t. We restrict our attention to the *transmission* delay caused by the limited bitrate that can be sent over a wireless channel (i.e., queuing delays of the packets at the transmitters are ignored). Furthermore, we consider only a single message, assuming that other messages (between other transmitters and receivers) are transmitted on orthogonal channels; therefore, interference does not play a role.

The power gain (inverse of the propagation attenuation) along the i-th edge is denoted as γ_i; its probability density function (PDF) is written as $f(\gamma_i)$. The PDF of the link gains are assumed to be independent. To make the following discussion more concrete we assume henceforth that the links undergo Rayleigh fading, i.e., the PDF of γ_i is [3]

$$f_{\gamma_{ij}}(\gamma_i) = \frac{1}{\overline{\gamma}_i} \exp\left[-\gamma_i/\overline{\gamma}_i\right], \ \ \gamma_{\mathrm{i}} \geq 0 \qquad (1)$$

where $\overline{\gamma}_i$ is the mean channel gain. The mean channel gains change very slowly, even with highly mobile nodes (typically, the means change $2-3$ orders of magnitude slower than the instantaneous link gains [3]. Thus information about the mean channel gains can be assumed to be available throughout the network.

[2]note again that routing based on geometrical information is *less general* than routing based on average CSI; the average channel gain includes not only the effect of the pathloss (which is related to the geometrical position of the devices) but also shadowing, random variations of pathloss coefficients, etc.

Information is communicated throughout the network by multiple-hop relaying with ideal physical-layer transmission. In other words, on each link transmission is done at link capacity, so that the transmission time for a message with source entropy H_{target} on link i

$$x_i = \frac{H_{\text{target}}}{\log\left[1+\gamma_i\right]}, \quad \text{for } \gamma_{\mathrm{i}} \geq 0 \qquad (2)$$

Since the links are Rayleigh fading, the PDF of the transmission delay over one link is [11]

$$f_{X_i}(x_i) = \frac{H_{\text{target}}}{\overline{\gamma}_i x_i^2} \exp\left[\frac{1}{\overline{\gamma}_i} + \frac{H_{\text{target}}}{x_i} - \frac{e^{H_{\text{target}}/x_i}}{\overline{\gamma}_i}\right] \qquad (3)$$

Note that this distribution has both an infinite mean and an infinite variance. It is also *subexponential*, meaning that it is more heavy-tailed than any exponential distribution. Subexponentials were intensely studied in the insurance literature [12] in the 1970s and 1980s, when catastrophic claims were sinking portfolios that appeared to be properly risk-balanced and re-insured. A key property is that the sum of iid subexponential variables is likely to be dominated by a single variable, thus any sample is likely to have extremely large values. Subexponential variables have several other properties that thwart standard methods of probabilistic inference and risk management, and also create special problems for finding routing with stochastic guarantees. It is a remarkable property of the algorithm developed in Sec. III that it works even for these extremely difficult distributions.

There are a number of ways how transmission at link capacity can be approximately achieved. If the instantaneous CSI is known at the transmitting node,[3] a (near) capacity achieving code, e.g., turbo-code or LDPC code, suitable for the specific SNR at hand, can be used. If the instantaneous CSI is not known, rateless codes [13], [14] can be used.

Our task is now to find a route such that maximizes the percentage of all channel states in which the route delay is no larger than a threshold t.

III. THEORY AND ROUTING ALGORITHM

In this section, we develop the mathematical framework for optimizing the probability of an event involving multiple random variables, particularly when integration is infeasible, which makes it impossible to reason about convolutions or moments. The chief result is a distribution-independent lower bound. We design a stochastic routing algorithm around this bound, give upper and lower bounds for the probability of on-time delivery, and show that the algorithm maximizes the lower bound. The results are very general and their application to statistics of the form Eq. (3), as done in Sec. IV, is only an example.

A. Mathematical preliminaries

Let X_i be a random variable; x_i be a realization of X_i; and, for a set of random variables $X_1, X_2, \cdots, X_n$, let E :

[3]Note that this requires only $local$ knowledge of instantaneous CSI; there is no need for network-wide knowledge of instantaneous CSI.

$x_1, \cdots, x_n \to \{0,1\}$ be an inequality that defines an event of interest as a polytope in $\mathbb{R}^n$. We will concentrate on the event of on-time delivery, written $E : \sum_i x_i \leq t$ for some deadline $t > 0$, and assume that each X_i has nonvanishing support on a continuous subset of the nonnegative reals, so that E is a closed set on $\mathbb{R}^n \geq 0$; these conditions are fulfilled for the distribution Eq. (3) as well as for many other practically relevant distributions.

Consider the nonlinear *probability map* to the unit hypercube $P : \mathbb{R}^n \to [0,1]^n$ defined by taking any realization $(x_1, x_2, \cdots, x_i)$ to the vector of probabilities $(F_1(x_1), F_2(x_2), \cdots, F_n(x_n))$, where $F_i(x) \doteq Pr(X_i \leq x_i)$. Applying map P to the event polytope E yields a hypercube subregion $P(E)$ whose boundary $P(\partial E)$ is typically curved. The significance of the map P is that probability is uniform in the hypercube, therefore the content (=hypervolume) of $P(E)$ is precisely the probability of event E, that is, $Pr(E) = \text{vol } P(E)$ is the quantity we are optimizing. For most distributions, we cannot evaluate the integral giving this volume, thus we seek distribution-independent bounds on $Pr(E)$. To that end we study the following property:

Definition 1: When $P(E)$ is a convex set, E is a *convex event* w.r.t. variables $X_1, X_2, \cdots$.

$P(E)$ is convex if (but not only if) E is convex in $\mathbb{R}^n$ and each CDF (cumulative distribution function) $Pr(X_i \leq t), t \geq 0$ is concave. A concave CDF implies a nonincreasing probability density function (PDF); this may be too restrictive. We will begin with concave CDFs but ultimately develop bounds for a much broader class of densities—those having nonincreasing right tails.

Knowing only that E is a convex event and the location of a point on the boundary $P(\partial E)$, we construct a bound by tightly fitting a diamond-shaped polytope inside $P(E)$. Figure 1 visualizes the bound for a two-variable event. Let us initially assume that each CDF has support on $(0, t] \in \mathbb{R}^+$, so that the boundary makes contact with each hypercube vertex $e_1, \cdots, e_n$ that adjoins the origin. Choose some point $\mathbf{p} \in P(\partial E)$ on the boundary; initially let us take $\mathbf{p} = p\mathbf{1}$, the point where the ray $\mathbf{1} = (1, 1, \cdots, 1)$ from the origin meets the boundary.

Lemma 1: (SIMPLE DIAMOND BOUND) $Pr(E) \geq \frac{1}{d!}(1 + q\sqrt{d})$ where $d \leq n$ is the number of random variables participating the event and $q = pd^{1/2} - 1$.

Proof: see Appendix VI-A. ■

The diamond polytope consists of two simplices spanning the points $\{\mathbf{0}, e_1, \cdots, e_n, \mathbf{p}\}$ and conjoined at a shared subsimplex spanning $\{e_1, \cdots, e_n\}$. In practice, $P(E)$ may not reach the hypercube corners $e_1, e_2, \cdots$ because for some variables, $Pr(X_i > t) > 0$. E.g., if the channel gain on a specific link is too low, then transmission over this single link already exceeds the admissible delay time. Let $m_i \doteq \max_x Pr(X_i \leq x|E)$ be the CDF value of the largest realization of X_i allowed by event E. The lower bound for the probability of on-time delivery is generalized as follows:

Lemma 2: (DIAMOND BOUND, CONVEX DISTRIBUTIONS) $Pr(E) \geq \frac{\prod_i m_i}{d!}(1 + q\sqrt{\sum_i m_i^{-2}})$ where $d \leq n$ is the number of random variables participating the event; $q = \langle (\mathbf{p}, 1), (\mathbf{z}, -1) \rangle / (\|\mathbf{p}\| \cdot \|\mathbf{z}\|)$ with $z_i = m_i^{-1}$; and any $\mathbf{p} \in P(\partial E) \subset [0,1]^n$.

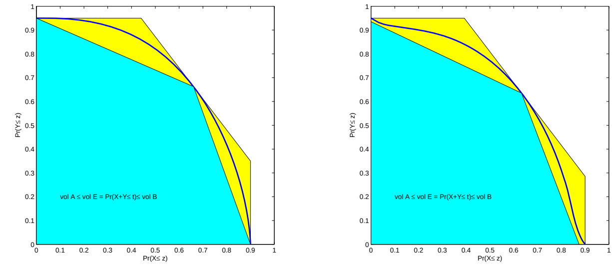

Fig. 1. A schematic of lower and upper diamond bounds on two-variable events. The area under the curve is the probability of the event. At left the event is convex; at right, near-convex, for which the bounds are adjusted.

Proof: See appendix VI-B. ■

In many routing problems the natural distribution functions do not yield a convex event. We extend the lower bound to such events by identifying a convex subvolume of $P(E)$.

Theorem 1: (DIAMOND BOUND, GENERAL DISTRIBUTIONS) Let $c_j = 0$ for concave F_j; otherwise $c_j = \max_{x|F_j''(x)=0, Pr(X_j \leq x|E)>0} x$, the largest sample value where the CDF of X_j inflects and E is feasible. If E is convex in $\mathbb{R}^n$, each density function $f_i(x) = F_i'(x)$ is nonincreasing on the right, and $\forall_i p_i \geq c_i$, then the convex bound holds with m_i set to the probability of the largest feasible value of X_i that satisfies E when $X_j = c_j, X_k = c_k, \cdots$. E.g., for $E : \sum_i x_i \leq t$, $m_i = \max Pr(X_i + (\max_{j \neq i} c_j) \leq t)$.

Proof: See Appendix VI-C. ■

B. Routing algorithm

In stochastic routing on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we have a combinatorial number of paths; for each source-target path $\mathcal{P} \subseteq \mathcal{E}$ we are interested in the event of on-time arrival, which we write $E|\mathcal{P} : \sum_{i \in \mathcal{P}} X_i \leq t$. Our goal is to find the path $\mathcal{P}$ maximizing the probability of this event $Pr(E|\mathcal{P}) = \text{vol } P(E|\mathcal{P})$. For most distributions this problem is NP-hard and sometimes inapproximable[15][16]; for the PDF in Eq. (3) it can even be challenging to numerically approximate vol $P(E|\mathcal{P})$ for a single path $\mathcal{P}$. However, with modest computation we can find a path that maximizes the lower bound given above. To do so, we search along a vector $\mathbf{v} \in [0,1]^{|\mathcal{E}|}$ for a point $\mathbf{p}$ on the boundary of the union of all events, $P(\partial(\bigcup_{\mathcal{Q} \in \text{st-paths}(\mathcal{G})} E|\mathcal{Q}))$:

1) Choose a bisection point $\mathbf{p}$ along vector $\mathbf{v}$.
2) For each edge random variable X_i, calculate the sample value $x_i = F_i^{-1}(p_i)$ that satisfies $Pr(X_i \leq x_i) = p_i$.
3) Find the shortest path $\mathcal{P}$ on $\mathcal{G}$ w.r.t. $x_1, x_2, \cdots, x_n$.
4) If $\sum_{i \in \mathcal{P}} x_i > t + \epsilon$, repeat bisecting closer to $\mathbf{0}$; if $\sum_{i \in \mathcal{P}} x_i < t - \epsilon$, repeat bisecting further from $\mathbf{0}$.

This bisection search terminates after no more than $\log 1/\epsilon$ instances of Dijkstra shortest path with a path whose realization lies on the boundary of E with precision $\epsilon \geq 0$.

C. Properties of the identified route

If the search vector is $\mathbf{v} = \mathbf{1}$, then the selected path $\mathcal{P}$ at $\mathbf{p} = p\mathbf{1}$ is robust to the random resampling of any single edge length in the following sense:

Proposition 1: Under single-edge resampling, the selected path is more likely to be in E than any other path.

Proof: See Appendix VI-D ■

No such guarantee is possible if we re-draw two or more edge lengths, because even though $Pr(X_i \leq x_i) = p$ and $Pr(X_j \leq x_j) = p$, it is possible that $Pr(X_i + X_j \leq x_i + x_j) \gg p$ due to nonlinearity of the distribution functions. The reader may intuit that this phenomenum is likely to favor the selected path more than any other path, and we support this intuition with the following result.

Theorem 2: The path selected at $\mathbf{p}$ maximizes the convex lower bound on probability of on-time arrival.

Proof: See Appendix VI-E ■

D. Remarks

The algorithm uses quantiles as a proxy for the delay distributions, and searches for the best set of quantiles for the routing-under-a-deadline task. The following discussion gives some mathematical intuition why *equiprobable* quantiles ($\mathbf{v} = \mathbf{1}$) are the most informative:

Definition 2: The *event shadow* $s(E, \mathcal{Q}, \mathbf{q})$ of path $\mathcal{Q}$ at point $\mathbf{q}$ is the set of all points $\mathbf{q}' \in P(E|\mathcal{Q})$ with $q_i \geq q'_i$.

The diamond bound is itself computed at $\mathbf{p}$ for $\mathcal{P}$ and in the shadow of $\mathbf{p}$ for all other paths. It is possible that the diamond bound could favor some other path $\mathcal{Q}$ if it were computed at some other point $\mathbf{r}$ on $\mathcal{Q}$'s event envelope that is outside the shadow of $\mathbf{p}$ (i.e., $\mathbf{r} \in P(\partial E|\mathcal{Q}) \backslash s(E, \mathcal{Q}, \mathbf{q})$). We argue that this outcome grows increasingly unlikely with graph size: The convexity of E implies that this point $\mathbf{r}$ must have some ordinates $r_i, r_j, \cdots$ that are substantially smaller than p, and therefore that much of the probability mass of $P(E|\mathcal{Q})$ is associated with unusually lucky draws from some of the edges in $\mathcal{Q}$. E.g., the threshold behavior of path $\mathcal{Q}$ is dominated by some edges with unusually broad distribution functions; lucky draws on these edges makes path $\mathcal{Q}$ suitable regardless of outcomes on its other edges, which in turn must have unusually narrow distribution functions. I.e., the path has mostly atypical delay distributions. However, if we have a smooth unimodal prior probability on the shape of delay distributions, the probability that such a path exists decays rapidly as we consider larger graphs.

E. Upper bounds

Here we sketch the construction of an upper bound on the probability of a convex event: A trivial hyperrectangle upper bound $Pr(E) \leq \prod_i m_i$ arises from the observation $\forall_i 0 \leq Pr(X_i < t) \leq m_i$. This bound can be sharpened shaving the far corner of the hyperrectangle with a cut through $\mathbf{p} \in P(\partial E)$ along the tangent space spanned by the derivatives of $P(\partial E)$ at $\mathbf{p}$. For our event, the $d-1$ vectors needed to determine the span can be calculated as $\frac{d}{dp_1} F_j(x_j - (F_1^{-1}(p_1) - x_1)) = \frac{-1}{f_1(F_1^{-1}(p_1))} f_j(x_j - (F_1^{-1}(p_1) - x_1)) = -f_j(x_j)/f_1(x_1)$ with $f_j(x_j) = F'_j(x_j)$ being the PDF and CDF of X_j at x_j. The cut volume and upper bound then follow from simple linear algebra.

We offer the following informal argument why the path $\mathcal{P}$ selected by our algorithm can also be expected to maximize this upper bound: For any alternate path $\mathcal{Q}$ and a location $\mathbf{q}$ in the event shadow of $\mathbf{p}$, we make two observations about the cutting hyperplane in the bound:

1) Since $\mathbf{q}$ is closer to $\mathbf{0}$ than $\mathbf{p}$, all else being equal, the hyperplane through $\mathbf{q}$ will cut off a larger volume.
2) Since some of the CDF values are reduced at $\mathbf{q}$, by concavity of CDFs the corresponding PDF values are increased and thus, *ceteris paribus*, the derivatives of $P(\partial E|\mathcal{Q})$ at $\mathbf{q}$ are more widely dispersed in value. This makes tangent space of $P(\partial E)$ at $\mathbf{q}$ less orthogonal to $\mathbf{1}$, which also increases the cut volume.

IV. Simulation results

For experimental validation, we compared the output of our algorithm against exhaustive search (where possible), and against a simple-minded approach, namely shortest-path for average CSI. One natural choice, average transmission time, is infeasible because our subexponentially distributed delay distributions have infinite means. Another natural choice of CSI cost statistic would be mean propagation attenuation, however we know this yields poor paths because it under-penalizes unreliable links. Median transmission time is a more competitive shortest-path statistic; it coincides with the first step of our algorithm and therefore allows us to assess the utility of finding the event boundary.

We performed several thousand trials comparing our algorithm against shortest path in a network with random placement of nodes. The transmitter is located at (0,0), and the receiver at (1,1). A set of 12 random nodes was placed at random in the unit square $[0, 1] \times [0, 1]$. The pathloss between two nodes was taken to be proportional to the squared distance between the nodes. Deadlines were chosen to be twice the transmission time that occurs when all channel gains attain their median value. In each trial we computed a route by our algorithm and by deterministic shortest path on the median transmission times, then sampled both paths 1000 times to estimate their probability of on-time delivery. [4] The chosen paths differed in $> 87\%$ of trials; of these the path chosen by our algorithm provided better on-time probability $> 94\%$ of the time. This is indicated by the circular dots massed above the diagonal in the probability scatter-plot in Fig. 2. In parallel experiments, the min-medians path is in turn almost always more successful than the min-squared-distances path; this is indicated by the square dots massed below the diagonal.

Results from further simulations with different settings (not shown here due to space restrictions) indicate the following: (i) for very small networks, where we were able to exhaustively enumerate and sample all paths to find the path with optimal on-time probability, our new algorithm found the true optimal path in most cases, and found only slightly suboptimal paths otherwise; (ii) with more stringent deadlines or low SNR links, our algorithm becomes increasingly dominant; and (iii) for networks with high SNR, a large number of hops, or very generous deadlines, the performance of deterministic routing on medians approaches that of stochastic routing.

[4] It must be noted that due to the sub-exponential behavior of the probability density function, 1000 samples may not always be sufficient to establish which of two paths is superior.

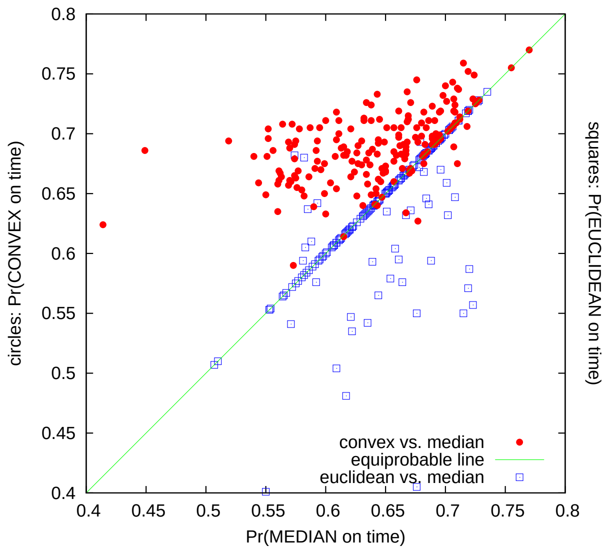

Fig. 2. Scatter plot comparing the probability of on-time delivery of the route chosen by our algorithm (vertical, red circles) versus that of the best path found by shortest path on medians (horizontal). Blue squares show medians are in turn better than squared distances.

V. Summary and Conclusions

This paper presented a new way of routing with delay guarantees in wireless ad-hoc networks. Finding a middle ground between flooding (which does not require the exchange of CSI, but is extremely energy-inefficient) and optimum routing based on instantaneous CSI in the network (which might require a large overhead for route discovery and maintaince of routing tables), we propose the use of *average* CSI, which has to be updated only very rarely. Even with this reduced CSI, it is possible to provide stochastic delay guarantees, i.e., to ensure that messages are delivered on-time in a percentage of cases given by the lower bound. We developed a novel, simple, yet highly effective algorithm to identify the route that most often fulfills the delay requirement. This algorithm is not heuristic, but rather based on analytical proofs for lower bounds on the success probability.

A possible alternative approach would be to allow only connections between nodes such that the mean channel strength between any pair of "connected" nodes exceeds the minimum required for a packet of this known length to be received with some chosen low probability of error for a chosen modulation. Then this same modulation can be used on every hop and the 'time' requirement becomes a 'number of hops' requirement, which has been studied in the literature [17]. However, this approach poses too stringent requirements on each possible link; there are thus many situations where our algorithm (where one fast link can compensate for the delay of another, slow, link) can find a route that fulfills the delay guarantee while the 'number of hops' algorithm fails.

The discussion of the algorithm concentrated on finding routes that give delay guarantees in Rayleigh fading channels. As a matter of fact, the algorithm itself is much more general and can be used in a variety of other applications. First and foremost, it is valid for *any* fading distribution, like Rice, Nakagami, etc. Since different fading distributions can occur in practical sensor networks [18], this easy generalizability is important. The method can also be used if the CSI is not the true average, but just some noisy or outdated estimate—just as long as the cumulative distribution function is known. Similarly, while we used the per-link transmission delay of an ideally coded system in our examples, the routing algorithm is not dependent on this assumption.

Furthermore, it is not necessary to restrict the QoS requirement to transmission delay. *Any* convex (e.g., additive) QoS constraints can form the basis of the algorithm. Last but not least, the restriction that the PDFs of the edge costs have to be independent can be lifted; details of this refinement will be reported in a future paper.

Acknowledgment

We thank Dr. Gulam Bhatti for valuable discussions, and Dr. Jin Zhang, Dr. Kent Wittenburg, and Dr. Joseph Katz for their support and encouragement.

VI. Appendix

A. *Proof of Lemma 1 in section III-A*

Because $P(E)$ is convex, it contains the convex hull of $\{\mathbf{0}, e_1, \cdots, e_i, \mathbf{p}\}$. This hull dissects into a standard simplex on $\{\mathbf{0}, e_1, \cdots, e_n\}$ and a regular simplex on $\{e_1, \cdots, e_n, \mathbf{p}\}$ that has been squashed along the ray $\mathbf{1}$. The ray intersects their common facet at $d^{-1/2}\mathbf{1}$ therefore the squashed simplex has height $q = pd^{1/2} - 1$. The content of the standard simplex is $1/d!$; the common facet is a regular simplex of $d-1$ dimensions with edge length $\sqrt{2}$, therefore its content is $\sqrt{2^{d-1}} \cdot \sqrt{d}/((d-1)!\sqrt{2^{d-1}}) = d\sqrt{d}/d!$. Extending this pyramidally to height q increases the content by factor q/d. Summing the contents give the result.

Remark 1: Consider the $[0,1]^{n-1}$ axis-aligned subspace containing event $P(E|X_k = 0)$. If we compute the point $\mathbf{q}_k$ where $\mathbf{1}$ meets this curve, then the content inside convex hull of $[e_1, \cdots, e_{k-1}, \mathbf{q}_k, e_{k+1}, \cdots, e_n, \mathbf{p}]$ lies inside $P(E)$ but outside the bound given above, and thus can be added to the bound to tighten it. We can do so holding each X_k=0, then also add the content in the convex hull of $\{e_1, \cdots e_{j-1}, \mathbf{q}_j, e_{j+1}, \cdots, e_{k-1}, \mathbf{q}_k, e_{k+1}, \cdots, e_n, \mathbf{p}\}$, etc.

B. *Proof of Lemma 2 in section III-A*

Using the Cayley-Menger determinant, the content of the lower simplex is $\prod_i m_i/d!$ and of the shared simplex is $(\prod_i m_i) d\sqrt{\sum_i m_i^{-2}}/d!$. The formula for q is the orthogonal distance from the shared simplex to any $\mathbf{p}$ (not just $\mathbf{p} = p\mathbf{1}$).

C. *Proof of Theorem 1 in section III-A*

Consider any two-dimensional slice through $P(E)$ and the axis e_i, viewed with e_i as the vertical axis. Because E is convex, the curve generated by the slice through the boundary $P(\partial E)$ is nonincreasing. If $\mathbf{p}$ can be located on this curve, then because $\forall_i p_i \geq c_i$, the curve has a central segment generated by the right tails of distributions, which must be convex. Project this segment onto e_i. By construction, m_i lies at or below the high end of the projection. Because the curve segment is nonincreasing and convex, any line drawn from $(0, m_i)$ to $\mathbf{p}$ lies wholly in $P(E)$. Thus any (upper) simplex with vertices $(m_1 e_1, \cdots, m_n e_n, \mathbf{p})$ lies in $P(E)$. By symmetry of argument, if the slice also passes through e_j, the line from $(0, m_i)$ to $(m_j, 0)$ is also in $P(E)$, implying the lower simplex is in $P(E)$.

D. Proof of Proposition 1 in section III-B

By construction, for any edge in the selected path, we have probability p of remaining in E. For any other path, a new draw must shorten the realized path length, thus the probability of entering E is $< p$.

E. Proof of Theorem 2 in section III-B

Any alternative path $\mathcal{Q}$ enters E by reducing some nonempty subset of its edge lengths $x_i, x_j, \cdots$, and thereby reducing the probabilities $Pr(X_i \leq x_i), Pr(X_j \leq x_j), \cdots$. Let $\mathbf{q}$ be a vector $(Pr(X_i \leq x_i), Pr(X_j \leq x_j), \cdots)$ of the probabilities of $\mathcal{Q}$'s realized edge lengths and let vector $\mathbf{p}'$ contain the corresponding values in $\mathbf{p}$. Recall that in the diamond bound, the content of the regular simplex is determined by the orthogonal distance of the sample point to the shared facet between the simplices. Since $\mathbf{q}^\top \mathbf{p}' < \mathbf{p}'^\top \mathbf{p}'$, this distance is reduced and the regular simplex is more squashed for $\mathcal{Q}$ than for $\mathcal{P}$, while all other elements of the bound are conserved.

References

[1] D. D. Perkins and H. D. Hughes, “A survey on quality-of-service support for mobile ad hoc networks,” *Wireless Communications and Mobile Computing*, vol. 2, pp. 503–513, 2002.

[2] T. Reddy, I. Karthigeyan, B. Manoj, and C. Murthy, “Quality of service provisioning in ad hoc wireless networks: a survey of issues and solutions,” *Ad Hoc Networks*, vol. 4, pp. 83–124, 2006.

[3] A. F. Molisch, *Wireless communications*. IEEE Press - Wiley, 2005.

[4] L. Georgiadis, M. J. Neely, and L. Tassiulas, “Resource allocation and cross-layer control in wireless networks,” *Foundations and Trends in Networking*, vol. 1, 2006.

[5] K. Akkaya and M. Younis, “A survey on routing protocols for wireless sensor networks,” *Ad Hoc Networks*, vol. 3, pp. 325–349, 2005.

[6] L. Liu and G. Feng, “Simulated annealing based multi-constrained qos routing in mobile ad hoc networks,” *Wireless Personal Communications*, vol. 41, pp. 393–405, 2007.

[7] D. Nguyen and P. Minet, “Qos support and olsr routing in a mobile ad hoc network,” in *Proc. ICN/ICONS/MCL 2006*, 2006.

[8] S. Draper, L. Liu, A. F. Molisch, and J. Yedidia, “Routing in cooperative wireless networks with mutual-information accumulation,” in *Proc. IEEE ICC 2008,*, p. submitted, 2008.

[9] T. Zhang, Y.; Gulliver, “Quality of service for ad hoc on-demand distance vector routing,” in *Proc. Int. Conf. Wireless And Mobile Computing, Networking And Communications 2005*, pp. 192–196, 2005.

[10] J. Novatnack, L. Greenwald, and H. Arora, “Evaluating ad hoc routing protocols with respect to quality of service,” in *Proc. Int. Conf. Wireless And Mobile Computing, Networking And Communications 2005*, pp. 205–212, 2005.

[11] A. F. Molisch, N. B. Mehta, J. S. Yedidia, and J. Zhang, “Cooperative relay networks using fountain codes,” in *Proc. IEEE Globecom* 2006.

[12] C. M. Goldie and C. Kluppelberg, “Subexponential Distributions,” In *A Practical Guide to Heavy Tails: Statistical Techniques for Analysing Heavy Tails* (R. Adler, R. Feldman and M.S. Taqqu, eds.) Birkhauser, Basel." 2007.

[13] M.Luby, “LT codes,” in *Proc. 43rd Annual IEEE Symp. Foundations of Computer Science(FOCS)*, pp. 271–282, 2002.

[14] A. Shokrollahi, “Raptor codes,” in *Proc. Int. Symp. Information Theory*, p. 36, 2004.

[15] E. Nikolova, M. Brand, D. Karger, “Optimal Route Planning under Uncertainty,” In *Proc., Int. Conf. on Automated Planning and Scheduling*, Lake District, England, 2006.

[16] E. Nikolova, J. A. Kelner, M. Brand, M. Mitzenmacher, “Stochastic Shortest Paths Via Quasi-convex Maximization,” In *Proc., European Symposium on Algorithms*, Springer LNCS 4168 pp. 552Â563. Zurich, 2006.

[17] S. Mukherjee and D. Avidor, “Connectivity, power and energy in a multihop cellular packet system,” In *Proc., IEEE Vehicular Techn. Conf. Fall*, pp. 1702-1707, 2005.

[18] S. Wyne, T. Santos, F. Tufvesson, and A. F. Molisch, “Channel measurements of an indoor office scenario for wireless sensor applications,” in *Proc. IEEE Globecom 2007*, 2007.