MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Delay-Energy Tradeoffs In Wireless Ad-Hoc Networks With Partial Channel State Information

Matthew Brand, Andreas Molisch

TR2008-078 December 2008

Abstract

Given a wireless network where each link undergoes small-scale (Rayleigh) fading, we consider the problem of routing a message from a source node to a target node while minimizing energy or power expenditure given a fixed time budget, or vice versa. Given instantaneous channel state information, we develop tight hyperbolic bounds on the quantities of interest and solve the related optimizations in closed form or via lightweight computations. If only average channel state information is available, probabilistical performance measures must be introduced. We therefore develop another set of bounds that supports resource-optimal routing with a guaranteed success probability. Our results rest on novel formulations and solution methods for hyperbolic convex programs and, more generally, nonlinear multicriterion combinatorial optimization.

IEEE Globecom 2008

Delay-energy tradeoffs in wireless ad-hoc networks with partial channel state information

Matthew Brand and Andreas F. Molisch *Fellow, IEEE*
Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA
email: {brand, molisch}@merl.com

***Abstract*—Given a wireless network where each link undergoes small-scale (Rayleigh) fading, we consider the problem of routing a message from a source node to a target node while minimizing energy or power expenditure given a fixed time budget, or vice versa. Given instantaneous channel state information, we develop tight hyperbolic bounds on the quantities of interest and solve the related optimizations in closed form or via lightweight computations. If only average channel state information is available, probabilistical performance measures must be introduced. We therefore develop another set of bounds that supports resource-optimal routing with a guaranteed success probability. Our results rest on novel formulations and solution methods for hyperbolic convex programs and, more generally, nonlinear multicriterion combinatorial optimization.**

***Index Terms*—stochastic routing, convex combinatorial optimization, wireless channel, Rayleigh fading**

I. Introduction

Wireless ad-hoc networks have emerged in recent years as an extremely important research area [1], [2]. In contrast to cellular communications, they avoid infrastructure, instead deploying a large number of low-cost nodes that forward the information. This approach not only decreases cost, but also decreases sensitivity to failure of a single link, e.g., due to bad propagation conditions. For these reasons, ad-hoc networks are promising, among other things, for factory automation and related applications that require ultra-reliable communications links [3].

Ultra-reliable wireless networks face two contradictory requirements: on one hand, the energy consumption has to be low, since nodes are battery operated, and exhausting the battery leads to node failure. On the other hand, the probability for successful transmission of data should be very high—by which we mean that a packet of data is transmitted from a source to a destination *within a prescribed delay*. We are interested in choosing a route (i.e., a sequence of nodes that pass the message along until it reaches the destination) that fulfills the delay constraints while using little energy. A simple solution to this problem uses a physical-layer transmission with fixed packet size and coding rate, where a packet is transmitted successfully, with a fixed time expenditure, if the link is good enough. Fulfilling the delay constraint is then equivalent to limiting the number of hops—a problem that has been well explored in the literature.

However, this simple approach ignores the possibility of speeding up transmission by investing more energy. For a single link, the tradeoff between transmission time and energy is straightforward: according to Shannon's capacity equation, the possible data rate increases with the transmit power. However, for networks with multiple hops, the tradeoff becomes much more complicated: it involves the question of which route to choose, as well as how much energy should be expended in each of the hops. In the current paper, we deal with this joint routing—energy/delay tradeoff problem for various amounts of channel knowledge.

To make the problem more precise, consider the following formulation: we are dealing with unicast in a network of N nodes, each of which can transmit with variable power and—due to adaptive modulation and coding (AMC)—can trade off transmission power and transmission time; such AMC capability is widespread in modern wireless systems[1]. A transmission is only considered successful if the message gets from the source to the destination with a delay that is smaller than a bound B (e.g., signals controlling a machine have to arrive within a very short time). We wish to find the routing and per-hop energy assignment that minimizes the overall energy expenditure while at the same time enabling a probability of successful transmission of q, where q is typically in the range of $q \in [90, 99.999]\%$ [2]. We assume that only the *statistics* of the channel state information (CSI) is available for the routing: 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 [4]). 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 discovery process often takes longer than the admissible delay of the information.

The problem is thus well-defined and practically relevant, but extremely hard to solve: There are on the order of $N!$ possible routes in the network, and for each route, the transmit energies of the nodes have to be optimized *under probabilistic constraints*. To our knowledge, there is no paper in the literature tackling this issue. The seminal work of [5] considered delay constraints but only with respect to scheduling on a single link; related work like [6], [7] considered the energy/delay tradeoff, but again only on a single link. A number of papers (e.g., [8] considered joint routing and power control,

[1] This is different from a scenario with fixed packet size and coding rate, where a packet has to be dropped on the link if the attenuation on a link is too strong.

[2] It is well-known that due to the randomness of wireless channels, costs can rise to ∞ as the guarantee approaches 100%.

but under the assumption of instantaneous CSI, and without delay constraints. Ref. [9] investigates routing with probabilistic delay constraints, but assumes fixed transmit power for each node.

In the current paper, we develop an near-optimum solution method that proceeds in several steps, each step drawing on a number of novel and recently-developed mathematical constructions. In a first, preliminary step (Sec. III), we investigate the optimum energy-delay tradeoff under the assumption of instantaneous CSI. We find the optimum energy allocation on a given route by bounding the delay with hyperbolic functions of the transmit energy in the region of interest (Sec. IV). Finding the optimum route can be done by means of a new, efficient mathematical technique that performs combinatorial convex optimization (Sec. V). We then show that the problem of statistical CSI can be reduced to that of the deterministic CSI by means of appropriate bounding techniques (Sec. VI). We stress that these algorithm are not heuristic; they fulfill provable bounds on the suboptimality.

Due to space constraints, all results in this paper are stated without proof. A longer manuscript with all proofs and a number of generalizations is available upon request.

II. System model

A. Deterministic setting

From the Shannon capacity equation, the transmission time per nat on a 1Hz bandwidth AWGN channel at full link capacity is

$$t = \log(1 + p\gamma)^{-1} \tag{1}$$

seconds, where p is the power (normalized by the noise power) and γ is the channel power gain, or inverse attenuation. Transmission time scales linearly with nats/bandwidth so all formulæ herein are on a per-nat/Hz basis. The inverse,

$$p = (\exp(1/t) - 1)/\gamma\ , \tag{2}$$

gives the power needed for a desired transmission time, with energy expenditure

$$e = pt = t(\exp(1/t) - 1)/\gamma\ , \tag{3}$$

of no less than $1/\gamma$ (because $\lim_{t\to\infty} e = 1/\gamma$). Conversely, spending $e \geq 1/\gamma$ energy units yields a transmission time of

$$\begin{aligned} t &= (-1/\gamma e - W_{-1}(-1/\gamma e \exp -1/\gamma e))^{-1} \\ &\leqslant (\log \gamma e + \log(1 + \log \gamma e))^{-1} \end{aligned} \tag{4}$$

seconds, where $W_{-1}(\cdot)$ is the branch of the multivalued Lambert W function that maps $[-\exp -1, 0) \to [-1, -\infty)$. The inequality is exact at $\gamma e = 1$ and a good approximation for $\gamma e > 1$. The relationships in Eq. (4) may be new to the literature.

These functions are convex decreasing on the positive line, specifying deeply "elbowed" resource trade-offs over the practical operating range of wireless devices. For example, eqn. 1 has the power-series approximation $t \approx 1/p\gamma + 1/2$. If one operates strictly in the subranges that lie on either side of the elbow (typically, high-bandwidth or low-SNR), linear approximations are useful [10], [11]. Here we seek nonlinear solution for the entire trade-off curve that are optimal or boundedly suboptimal.

Given a relay or network of wireless channels, each with a unique channel gain γ_i and convex decreasing resource trade-off $y_i = f_i(x_i)$, we pose the following questions:

1) **Allocation**: Given a series of N links and a total (per-nat/Hz) budget B on resource x, what allocation minimizes total use of resource y?

$$\min \sum_i f_i(x_i) \quad \text{such that} \sum_i x_i \leq B \tag{5}$$

2) **Routing**: What route through a network $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ of $M = |\mathcal{E}|$ links affords the optimal allocation?
3) **Re-allocation**: Given partial or novel information, can the optimal allocation be revised on-the-fly?

These deterministic optimizations are useful for small paths and networks where channel state measurements remain valid long enough to be acted upon.

B. Stochastic setting

In many settings, instantaneous channel state information may not be measurable or constant over the time scale of interest, so we must work with the probability distribution over channel states. Instead of optimizing use of one resource subject to a constraint on another, we minimize use of one resource subject to a *bound on the probability of success* in meeting a constraint on the other, i.e. Eq. (5) is replaced with

$$\min \sum_i y_i \quad s.t.\text{Pr}(\sum_i X_i \leq B) \geq q. \tag{6}$$

Here y_i is a resource allocated to the i^{th} link, X_i is a random variable whose PDF is parameterized by y_i, B is a budget, and q is a minimal acceptable probability of success.

1) Distributions over time and power costs: Stochasticity arises in real settings because the channel gain γ is random variable that is exponentially distributed with mean $\overline{\gamma}$. Solving eqn. 2 for γ reveals that $(-1 + \exp 1/t)/p$ is exponentially distributed with parameter $\theta = 1/\overline{\gamma}$, while $p/(-1 + \exp 1/t)$ is inverse-gamma distributed with parameters $\alpha = 1, \beta = 1/\overline{\gamma}$. Solving these for the conditional time and power CDFs yields

$$\begin{aligned} Pr(t \leq x|p) &= \exp \frac{1 - \exp 1/x}{p\overline{\gamma}} \\ Pr(p \leq x|t) &= \Gamma(1, \frac{-1 + \exp 1/t}{x\overline{\gamma}}) = \exp \frac{1 - \exp 1/t}{x\overline{\gamma}} \end{aligned}$$

where $\Gamma(a, b) \doteq \int_b^\infty t^{a-1} e^{-t} dt$ is the incomplete gamma function. These CDFs give the probability of meeting a time (resp., power) constraint given an expenditure of power (resp., time). Similarly, energy and time have the stochastic trade-off

$$\begin{aligned} Pr(t \leq x|e) &= \exp\{x(1 - \exp 1/x)/(e\overline{\gamma})\} \\ Pr(e \leq x|t) &= \exp\{t(1 - \exp 1/t)/(x\overline{\gamma})\} \end{aligned}$$

These subexponential distributions have several unfavorable properties: They are more heavy-tailed than any distribution in the exponential family. They are not closed under convolution,

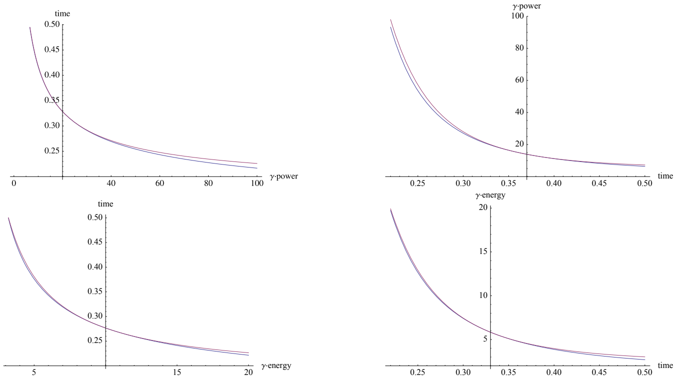

Fig. 1. Trade-off and upper bound curves for typical operating ranges of wireless devices. Vertical axes are placed where bounds are exact. Tighter bounds are available on these ranges and on subranges.

so the sum of random variables in eqn. 6 cannot be evaluated. They have infinite moments, so one cannot reason about expectations. In short, it is difficult if not impossible to analytically compare combinations of these distributions.

Even the deterministic setting is challenging: the allocation problem is a convex combination of convex functions, therefore in principle it can be solved numerically through various convex optimization techniques [12]. However the numerical effort can still be considerable and might require more time and energy than is being saved. To be useful, solutions should be computationally lightweight and sufficiently accurate to be reliable in resource-starved situations. To that end, we will introduce a family of tight upper bounds on the deterministic trade-offs, then solve the bounded allocation problems in closed form. This in turn yields a link characterization that supports near-optimal routing solutions in $O(M \log^2 M)$ time for networks of M links. Finally, we show how the stochastic problem can be transformed into a deterministic problem and solved.

III. Hyperbolic bounds

Our solution begins in the deterministic setting, by upper bounding the trade-off at each link over some finite practical range with a hyperbolic curve of the form $h_i/x_i^n + c_i$ for some global exponent $n > 0$ and a unique scale factor h_i and offset c_i for each link. Figure 1 shows that as approximations to the deterministic trade-offs, these are quite good over a large operating range; with a suitable choice of n the expected approximation error can always be driven down to a few percent. Indeed, in many cases we can analytically bound the maximum and expected error.

Specifically, we upper-bound a convex decreasing curve $y = f(x)$ with a hyperbolic curve $h/x^n + c$ on some interval $x \in [x_0, x_1]$, $y \in [y_1 = f(x_1), y_0 = f(x_0)]$ using one of four tactics:

(A) Set the two curves to meet at endpoints x_0, x_1 with

$$h = \frac{y_0 - y_1}{x_0^{-n} - x_1^{-n}}, \qquad c = \frac{y_1 x_1^n - x_0^n y_0}{x_1^n - x_0^n}, \tag{7}$$

choosing n small enough to guarantee an upper bound;

(B) make the curves tangent at some point $x = \mu \in [x_0, x_1]$ with

$$\begin{aligned} h &= -\mu^{n+1} f'(\mu)/n \\ c &= f(\mu) + \mu f'(\mu)/n \end{aligned} \tag{8}$$

choosing n large enough to guarantee an upper bound and μ to minimize the expected or maximum gap;

(C) fit h, c, n to give a good approximation of $f(x)$ and adjust the offset c to make an upper bound; or

(D) an upper bound on one trade-off can be inverted to yield a shifted hyperbolic upper bound on the inverse trade-off:

$$h/x^n + c \geq f(x) \iff \sqrt[n]{h}/\sqrt[n]{y - c} \geq f^{-1}(y) \tag{9}$$

As an example, for time t as a function of power p (Eq. (1)), the hyperbolic bound $t \leq h/p^n + c$ can be fit by making the curves tangent at some $\mu > 0$ as per Eq. (8):

$$h = \frac{\gamma \mu^{n+1}}{n(1 + \gamma\mu)\log(1 + \gamma\mu)^2}, \qquad c = \frac{1}{\log(1 + \gamma\mu)} - \frac{h}{\mu^n}.$$

Proposition 1: For all positive p, μ and $n \geq 1$, this hyperbolic curve is an upper bound, with equality at $p = \mu$.

Remark 1: The approximation can be tightened by choosing $0 < n < 1$, however the resulting hyperbolic may be an upper bound only in some finite interval around μ.

One may choose the point of tangency μ to minimize the expected approximation error with respect to some distribution on power costs. For example, consider the expected additive error, $\int h/p^n + c - 1/\log(1 + \gamma p)\, d\text{Pr}(p)$, where h, c are defined as above and $\text{Pr}(p)$ is an exponential density with mean λ. The integral does not converge, but its derivative w.r.t. μ has a single zero on $0 < \mu < \infty$ at $\mu = \lambda \sqrt[n]{(e^{-p_0/\lambda} - e^{-p_1/\lambda})/(\Gamma(1-n, p_0/\lambda) - \Gamma(1-n, p_1/\lambda))}$ (assuming some finite operating range $p \in [p_0, p_1]$). In a typical operating regime ($\gamma p \in [4, 100]$ SNR, distributed exponentially with mean $\lambda = 10$), the additive error is minimized at $\mu \approx 0.95\lambda$ at $n = 1$ and $\mu \approx 0.99\lambda$ at $n = 3/4$. Bounds on maximal and expected error follow algebraically. Less formally, by simply setting $\mu = \lambda$ and numerically calculating the expected multiplicative error, we find that the hyperbolic bound overestimates power costs by $< 2.4\%$ on average at $n = 1$ and $< 1.0\%$ at $n = 3/4$.

Similar results for all trade-offs are available in the extended manuscript.

IV. Optimal deterministic allocation

With hyperbolic upper bounds $h_i/x_i^n + c_i \geq f_i(x_i)$, the allocation problem takes the form

$$\min \sum_i h_i/x_i^n + c_i \text{ s.t. } \sum_i x_i \leq B, \quad \forall_i x_i > 0 \tag{10}$$

W.l.o.g., we drop the fixed cost $\sum_i c_i$ and upgrade the simplex constraint to equality at $\sum_i x_i = B$. This we can solve in closed form:

Theorem 1—Hyperbolic Programs: The problem

$$\min \sum_i h_i/x_i^n \text{ s.t. } \sum_i x_i^m = B, \quad \forall_i x_i > 0 \tag{11}$$

for $n > 0, m \geq 1, h_i > 0, x_i \geq 0$ has unique solution

$$x_i = \sqrt[m]{B \cdot h_i^{m/(m+n)} / \sum_k h_k^{m/(m+n)}} \tag{12}$$

with value $B^{-n/m} \| h_1, h_2, \cdots \|_{m/(m+n)}$ where $\| \cdot \|_p$ is the L_p quasi-norm.

Corollary 1: For the full allocation problem (Eq. (10)) the cost is upper bounded by

$$B^{-n/m} \| h_1, h_2, \cdots \|_{m/(m+n)} + \| c_1, c_2, \cdots \|_1 \tag{13}$$

For the remainder of this paper, we will assume $m = 1$.

A simple example: The optimal $n = 1$ allocation assigns to the ith link $x_i = B\sqrt{h_i} / \sum_j \sqrt{h_j}$ of the budgeted resource and $y_i \leq \sqrt{h_i} \sum_j \sqrt{h_j} / B + c_i$ of the minimized resource. It can be shown that for $n = 1$ hyperbolic bounds on the deterministic trade-offs, the scale factor varies with the channel attenuation: $h_i \approx z/\gamma_i$ for some constant z. Thus, for example, to minimize power use (say, to minimize RF interference with other wireless devices) and guarantee a delivery deadline, power and time should be approportioned according to the square root of the channel attenuation. The allocation inherits suboptimality bounds from the hyperbolic bounds, so in this case the expected suboptimality would be $< 2.4\%$.

Before moving on to re-allocation, we point out a closed-form allocation for horizontally shifted hyperbolics:

Corollary 2: $\min \sum_i h_i/(x_i + v_i)^n + c_i$ s.t. $\sum_i x_i^m = B$ is solved at $x_i = ((B + \sum_k v_k) h_i^{m/(m+n)} / \sum_k h_k^{m/(m+n)} - v_i)^{1/m}$.

The shift (v_i) allows more flexible fits to convex trade-off curves, but unshifted hyperbolics will prove much more versatile in the routing sequel.

A. On-the-fly and distributed allocation

After computing an allocation and transmitting along the path, one might reach the ith link and discover its channel gain has drifted. Revising the optimal resource allocation on-the-fly is a matter of simple algebra. To do so, it is useful to propagate and update the partial sum $S_i \doteq \sum_{j \geq i} h_j^{1/(1+n)}$ and the remaining time-to-deadline B_i. Then if the hyperbolic scale factor for link i changes from h_i to h_i', for $n = 1$ the minimal allocation at link i changes from $\sqrt{h_i} S_1 / B_1 + c_i$ to $(h_i' + \sqrt{h_i'} S_{i+1})/B_i + c_i$.

This immediately suggests a range of distributed algorithms where we do not determine each h_i from instantaneous measurements but instead assume a value for each h_i on the basis of some historical statistic, then update allocations on-the-fly as above. This idea is developed more precisely in the stochastic setting in section VI-C.

V. Near-optimal routing

It is possible to compute near-optimal routes with respect to hyperbolic bounds. To do so, we first introduce the idea of a *linear* multicriterion combinatorial optimization, where the objective is a weighted average of multiple criteria. Typically one wants to reason about the optimum *before* the weighting is known; to do so one considers the entire set of possible solutions, indexed by the weighting parameters. Here we will consider a bicriterion path cost $C(\mathcal{P}, \lambda) \doteq \sum_{k \in edges(\mathcal{P})} w_k(\lambda)$ where $\mathcal{P}$ is a path and λ is a weighting parameter that balances two criteria to determine each edge length:

$$w_k(\lambda) = h_k^{m/m+n} \lambda + B^{n/m} c_k \, .$$

This defines a Bicriterion Shortest Path (BSP) problem that maps each source-target path $\mathcal{P}_i$ in the network to a line in the positive quadrant ($\lambda \geq 0$) with slope $a_i = \sum_{k \in edges(\mathcal{P}_i)} h_k^{m/m+n}$, intercept $b_i = B^{n/m} \sum_{k \in edges(\mathcal{P}_i)}^{n/m} c_k$. The key property is that on each line $a_i \lambda + b_i$, there is a point at $\lambda = a_i^{n/m}$ that indexes the *nonlinear* cost of the corresponding path under the optimal allocation, i.e.,

$$a_i \lambda + b_i \propto B^{-n/m} \| h_1, h_2, \cdots \|_{m/(m+n)} + \sum_k c_i . \tag{14}$$

We use this geometry as a scaffolding to find the minimal hyperbolic-cost route. Figure 2 depicts the bundle of lines corresponding to all source-target paths, each ornamented with its cost point as per Eq. (14). The infimum of this bundle of lines is a piecewise linear concave curve we call the bundle boundary. The lines forming the infimum are the BSP solution set; each represents a shortest path on some λ-interval, and can be queried in $O(M \log M)$ time by Dijkstra's shortest path algorithm on scalar edge weights $w_i(\lambda)$ generated by an appropriate value of λ. Our solution revolves around a curve formed by rolling a line $a\lambda + b$ around the top of the boundary, pivoting on its vertices, and tracing the evolution of the point at $\lambda = a^{n/m}$.

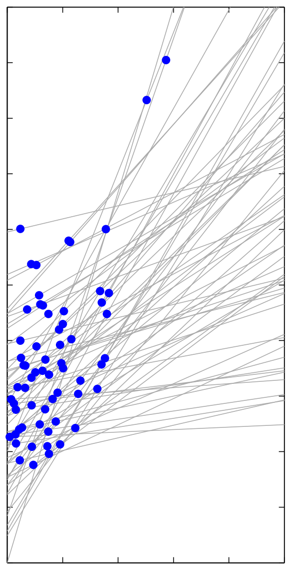

Fig. 2. BSP bundle with costs.

Theorem 2: For positive m, n, this curve is convex, piecewise smooth, and a lower envelope on all possible optimal cost points for any network having the same BSP boundary.

Figure 3 illustrates that this envelope connects the cost points of the boundary paths. Therefore one of these paths is either optimal or near-optimal with the following suboptimality bounds:

Theorem 3—Hyperbolic Min-Cost Routing: Let $\mathcal{P}_i, \mathcal{P}_j$ be two paths whose lines intersect on the bundle boundary at $\lambda_{ij} = \frac{b_j - b_i}{a_i - a_j}$ with $a_j < a_i$. The boundary has one such pair that contains either the optimal route or a near-optimal route with additive suboptimality upper-bounded by $\min_{ij}(a_i - \frac{1}{2}\lambda_{ij})^2$ for $n = m = 1$ and $\min_{ij} a_i(a_i^{n/m} - \lambda_{ij}) + \frac{n}{m}(\frac{m\lambda_{ij}}{m+n})^{m+n/n}$ for general m, n. A looser bound $((a_i^{n/m} - a_j^{n/m})/2)^{m+n/n}$ yields an $n = m = 1$ suboptimality ratio of $\frac{a_i^2 + b_i}{(3a_i - a_j)(a_i + a_j)/4 + b_i}$.

Remark 2: Note that the population of boundary lines grows at least polynomially with the number of edges, with $\lim_{|\mathcal{E}| \to \infty} a_i - a_j = 0$; thus suboptimality vanishes asymptotically with graph size.

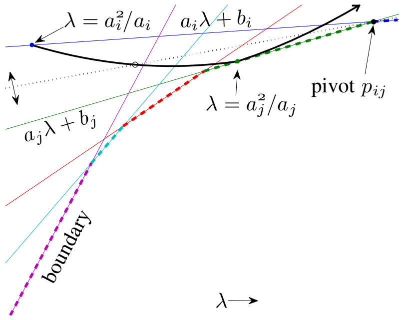

Fig. 3. Construction of the nonlinear cost envelope, quadratic case. Each index is associated with a line of the form $\ell_i(\lambda) \doteq a_i\lambda + b_i$ in the positive quadrant. The figure shows the subset of lines that form the piecewise linear bundle boundary $\inf_i \ell_i(\lambda)$ for all $\lambda \geq 0$. Each path's cost is a point on its line at coordinate $\lambda = a_i, \ell_i(a_i) = a_i^2 + b_i$. The cost envelope is a piecewise parabolic curve formed by pivoting a line on each intersection point on the boundary, and tracing the locus of points satisfying $\lambda = a^2/a = a$ as the slope a of the line varies between those two envelope lines meeting at the pivot point. The cost envelope and boundary cost points give a lower and upper bound on the mininum possible cost $\min_i a_i^2 + b_i$ attainable by any path in the graph.

Corollary 3: If $n = 1$ and $\forall_{ij}\, c_i = c_j$, the routing problem reduces to min-cost path on a graph with edge weights $w_i = \sqrt{h_i}$.

The BSP solution set can be explicitly enumerated and scanned for the best path, with polynomial smoothed time complexity. However, since the envelope can be characterized parametrically and differentiated, we can perform a bisection search on the slopes of the boundary lines to find the boundary path having lowest nonlinear cost, using envelope derivatives to decide the correct bisection interval.

Theorem 4: A path whose boundary line has slope a_i within ϵ of optimal can be found in $O(M \log^2 M \log 1/\epsilon)$ time.

VI. Stochastic Methods

As argued above, the subexponential nature of the distributions of the optimization variables precludes analytic approaches. We advocate lower-bounding the probability of success with a more tractable family of linearly additive phase-type distributions. There are several possibilities; in this paper we will use gamma distributions with a common spread parameter β and varied location parameters α_i; e.g., if $X_1 \sim \text{Gamma}(\alpha_1, \beta)$ and $X_2 \sim \text{Gamma}(\alpha_2, \beta)$, then $X_1 + X_2 \sim \text{Gamma}(\alpha_1 + \alpha_2, \beta)$.

For each link, we assume a finite operating range and choose a parameterized bounding distribution that is dominated by the true distribution on that interval, e.g., $\forall_{t_{\min} \leq t \leq t_{\max}} Pr_{bounding}(X_i \leq t|y_i, \alpha_i) \leq Pr_{true}(X_i \leq t|y_i)$. Any reasoning done with the bounding distributions then underestimates the true probability of success. For example, minimizing power subject to time constraints, we may conservatively choose $t_{\max}$ to be the full time budget and $t_{\min}$ to be the shortest single-link transmission time attainable at max power. That makes the set of successful events a subset of the set of events for which the bound is valid.

Because the true distribution is subexponential, there is some crossing point $t_c > 0$ below which the true distribution dominates. [3]. Setting the two CDFs equal at $t_c = t_{\max}$ guarantees the validity of the bound over the operating range while determining a functional relationship $\alpha_i = g_i(y_i)$ between resource use y_i and the parameter α_i of the bounding CDF. Generally, increasing resource use y_i decreases parameter α_i, which in turn increases the probability $Pr(X_i \leq t)$. Trivially, given any *fixed* resource allocation $y_1, y_2, y_3, \cdots$ over the whole network and linearly additive bounding distributions, the min-cost path on edge costs $\alpha_i = g_i(y_i)$ maximizes the lower bound on probability of success.

A. Resource allocation on a fixed route

Now consider the allocation problem on a path with stochastic resource trade-offs. The linkwise bounding distributions have been constructed so that for any path of links, there is a distribution $G_\beta(A, B) \doteq \text{Pr}_{X \sim \text{Gamma}(A,\beta)}(X \leq B)$, which, for $A \geq \sum_i \alpha_i$, lower-bounds the true probability of success on that path, i.e., $\text{Pr}_{true}(\sum_i X_i \leq B) \geq G_\beta(A, B)$ for $B \leq t_{\max}$ (at least). To obtain a specific probability of success q, we set $G_\beta(\sum_i \alpha_i, B) = q$ and invert G_β on its first parameter to obtain a new constraint $\sum_i \alpha_i \leq A = G_\beta^{-1}(q, B)$. Since resource use can also be expressed in terms of α_i as $y_i = g_i^{-1}(\alpha_i)$, our goal (Eq. (6)) can be rewritten

$$\begin{aligned} \min \sum_i g_i^{-1}(\alpha_i) \\ \text{s.t.} \sum_i \alpha_i \leq G_\beta^{-1}(q, B)\,. \end{aligned} \tag{15}$$

We may now employ the same hyperbolic bounding schemes and optimization methods developed in the deterministic setting to solve for the optimal α_i, and then calculate resource allocations $y_i = g_i(\alpha_i)$. We illustrate by working out the case of minimal power use, subject to a time constraint. To guarantee the lower bound on the probability of success, at each link we set the gamma and time CDFs equal at $t = t_{\max}$ and solve for the needed power p_i, yielding

$$p_i = g_i^{-1}(\alpha_i) = \frac{1 - \exp 1/t_{\max}}{\overline{\gamma}_i \log(\Gamma(\alpha_i, t_{\max}\beta)/\Gamma(\alpha_i))}.$$

A power series expansion about $\alpha_i = 0$ reveals the hyperbolic approximation

$$p_i = g_i^{-1}(\alpha_i) \leq \frac{h_i}{\alpha_i} + c_i for h_i = \overline{\gamma}_i^{-1} \frac{1 - \exp 1/t_{\max}}{\Gamma(0, t_{\max}\beta)} \tag{16}$$

which becomes an upper bound with suitable choice of c_i. With this Eq. (12) can be applied directly to Eq.(15) to compute the optimal α_i w.r.t. the hyperbolic bounds, which in turn gives the optimal allocation y_i w.r.t. the hyperbolic and gamma bounds jointly.

The same construction is used to minimize time given an power constraint, except with

$$\begin{aligned} t_i = g_i^{-1}(\alpha_i) &= \log(1 - \overline{\gamma}_i p_{\max} \log(1 - \Gamma(\alpha_i, p_{\max}\beta)/\Gamma(\alpha_i)))^{-1} \\ &\leq \frac{h_i}{\alpha_i} + c_i for h_i = \overline{\gamma}_i^{-1} \frac{1}{p_{\max}\Gamma(0, p_{\max}\beta)} \end{aligned}$$

[3]In principle, the bounding distribution can also dominate in a very small, low-probability interval between 0 and $t_{c'} \approx 0$. Typically $t_{c'} \ll t_{\min}$, so our bounds remain valid. In the case of gamma bounds, $t_{c'}$ can be driven toward 0 by increasing β.

For energy given time, $g^{-1}(\alpha)$ and h_i are the same as those of power given time, except multiplied by $t_{\max}$. For time given energy, let $\nu_i \doteq -\overline{\gamma}_i e_{\max} \log(1 - \Gamma(\alpha_i, e_{\max}\beta)/\Gamma(\alpha_i))$. Then

$$\begin{aligned} e_i = g_i^{-1}(\alpha_i) &= \{-1/\nu_i - W_{-1}(-1/\nu_i \exp -1/\nu_i)\}^{-1} \\ &\leq \{\log \nu_i + \log(1 + \log \nu_i)\}^{-1} \\ &\leq \frac{h_i}{\alpha_i} + c_i for h_i \propto \overline{\gamma}_i^{-1} \end{aligned} \tag{17}$$

however here the precise value of h_i must be determined using the methods in section III.

There is a pleasing symmetry with $h_i \propto \overline{\gamma}_i^{-1}$ in all cases under $n = 1$ hyperbolic bounds, therefore the optimal setting of the gamma location parameters is $\alpha_i \propto \overline{\gamma}_i^{-1/2}$. However, in implementation, there is an asymmetry: When minimizing power, one simply transmits at the allocated power; when minimizing time, the transmitter and receiver must handshake to learn the current channel gain, then one transmits using whatever power is necessary to make the allocated time.

If the requested probability of success is too high, the optimal allocation will put one or more links outside of their operating range. Resource use at these links can be clamped to their maximums and the allocation problem re-solved at the remaining links. If this fails to yield a viable allocation we report that the desired probability of success is infeasible w.r.t. the chosen bounds.

B. Routing

The routing machinery of section V can be used to find near-optimal paths w.r.t. the gamma and hyperbolic bounds.

C. On-the-fly reallocation

Uncertainty is reduced as a multi-hop relay progresses: The probability distribution narrows and only B_i of the budgeted resource remains once we have reached link i. Re-solving Eq. (15) at link i gives a revised $n = m = 1$ upper bound of $\sqrt{h_i} \sum_{j \geq i} \sqrt{h_j}/G^{-1}(B_i, q)$.

If link i's probability distribution is collapsed *prior* to transmission (e.g., by measuring its instantaneous channel gain), we learn the deterministic trade-off $y_i = f_i(x_i)$. Conditioning the probability of success on this information yields a modification of the optimization problem in Eq. (15):

$$\begin{aligned} &\min f_i(x_i) + \sum_{j>i} g_j^{-1}(\alpha_j) \\ &\text{s.t.} \sum_{j>i} \alpha_j \leq G^{-1}(B_i - x_i, q) \end{aligned} \tag{18}$$

The nonlinear dependence of the probability of success on x_i presents a difficulty. To solve with $n = 1$ bounds, we lower-bound the probability of success with an affine function $b - ax_i \leq G^{-1}(B_i - x_i, q)$, upper-bound the stochastic link costs $g_j^{-1}(\alpha_j)$ as above, and upper-bound the deterministic cost $f_i(x_i)$ with $h_i'/x_i + c_i'$ as in section IV-A, then solve the hyperbolic program

$$\begin{aligned} \min \quad & h_i'/x_i + c_i' + \sum_{j>i} h_j/\alpha_j + c_j \\ \text{s.t.} \quad & \alpha_i + \sum_{j>i} \alpha_j \leq b, \qquad \alpha_i = ax_i \end{aligned} \tag{19}$$

to obtain the minimal safe expenditure at link i of $y_i = f_i(ab\sqrt{h_i'}/(\sqrt{h_i'} + \sum_{j>i} \sqrt{h_j}))$. The reader may recognize this as a variation on the on-the-fly allocation algorithms sketched in section IV-A.

VII. Summary and Conclusions

In the wireless setting, power/energy can be traded-off with transmission time, due to channel capacity constraints, and it is desirable to optimize one resource while guaranteeing quality of service w.r.t. another. In this paper, we have developed a framework for optimizing such tradeoffs. More generally, the results are applicable to convex decreasing trade-offs in combinatorial settings such as finding the min-cost path w.r.t. a resource trade-off that is budget-constrained. Our main results are: tight hyperbolic bounds and a closed form resource allocation with less than 3% expected suboptimality; on-the-fly reallocations; a low-complexity combinatorial solution for finding a path with a near-optimal trade-off; bounds on its suboptimality; and the extension of all these methods to the stochastic setting where the trade-off curves are known only probabilistically. Suitably generalized, these methods may prove useful in any setting where there is a combinatorial optimization subject to stochastic resource trade-offs and budgetary constraints.

Acknowledgment

We thank Drs. Gulam Bhatti, Jin Zhang, Kent Wittenburg, Joseph Katz for their support, encouragement, and questions.

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] SP100, "Call for proposals - wireless for industrial process measurement and control," Tech. Rep. ISA/SP100.11-2006-001R8, ISA, 2006.

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

[5] R. A. Berry and R. G. Gallager, "Communication over fading channels with delay constraints.," *IEEE Transactions on Information Theory,*, vol. 48, pp. 1135–1149, 2002.

[6] X. Zhong and C.-Z. Xu, "Delay-constrained energy-efficient wireless packet scheduling with qos guarantees,"

[7] Y. Yang, V. Prasanna, and B. B. Krishnamachari, "Energy minimization for real-time data gathering in wireless sensor networks," in *IEEE Trans. Wireless Communications*, vol. 5, 2006.

[8] R. L. Cruz and A. V. Santhanam, "Optimal link scheduling and power control in cdma multihop wireless networks," in *Proc., IEEE Globecom*, 2002.

[9] M. Brand, P. Maymounkov, and A. Molisch, "Convex bounds for arrive-on-time stochastic routing," in *in reviews*, 2008.

[10] X. Liu and A. Goldsmith, "Optimal power allocation over fading channels with stringent delay constraints," in *Proc., IEEE ICC*, vol. 3, pp. 1413–1418, 2002.

[11] A. Fu, E. Modiano, and J. Tsitsiklis, "Optimal energy allocation for delay-constrained data transmission over a time-varying channel," in *Proc., IEEE INFOCOM*, vol. 2, pp. 1095–1105, 2002.

[12] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.