

Frugal Routing on Wireless Ad-Hoc Networks

Gunes Ercal, Rafit Izhak-Ratzin, Rupak Majumdar, and Adam Meyerson

University of California, Los Angeles

Abstract. We study game-theoretic mechanisms for routing in ad-hoc networks. Game-theoretic mechanisms capture the non-cooperative and selfish behavior of nodes in a resource-constrained environment. There have been some recent proposals to use incentive-based mechanisms (in particular, VCG) for routing in wireless ad-hoc networks, and some frugality bounds are known when the connectivity graph is essentially complete. We show frugality bounds for random geometric graphs, a well-known model for ad-hoc wireless connectivity. Our main result demonstrates that VCG-based routing in ad-hoc networks exhibits small frugality ratio (i.e., overpayment) with high probability. In addition, we study a more realistic generalization where sets of agents can form *communities* to maximize total profit. We also analyze the performance of VCG under such a community model and show similar bounds. While some recent truthful protocols for the traditional (individual) agent model have improved upon the frugality of VCG by selecting paths to minimize not only the cost but the overpayment, we show that extending such protocols to the community model requires solving NP-complete problems which are provably hard to approximate.

1 Introduction

We study the frugality ratio (FR), a measure of cost-efficiency, of the generalized VCG mechanism for reliable routing in the presence of non-cooperative behavior in ad-hoc networks. We model ad-hoc networks by random geometric graphs (RGG), and show that VCG-based routing exhibits small frugality ratio with high probability (w.h.p.). We generalize the standard model of agent behavior by allowing sets of nodes to form communities to maximize the total profit and demonstrate bounds on the frugality ratio for this model as well. Moreover, while some recent truthful protocols for the traditional (individual) agent model have improved upon the frugality of VCG by selecting paths to minimize not only the cost but the overpayment, we show that extending such protocols to the community model requires solving NP-complete problems which are provably hard to approximate.

Reliable and cost-efficient routing in ad-hoc networks is a well-studied problem, with numerous proposals for routing protocols. Many of these protocols assume that the nodes in the network behave co-operatively. In resource-scarce environments, such as ad-hoc networks, this co-operativeness assumption is suspect. Forwarding a packet incurs some cost and in the absence of other incentives, nodes belonging to one community may refuse to forward packets belonging to

another community. Under these assumptions, it is more reasonable to model a network as a game played between independent selfish agents, and to apply game theoretic reasoning to develop incentive-based routing protocols [1,2].

In an incentive-based routing protocol, a node is paid monetary compensation in return for forwarding a packet. The compensation covers the cost incurred by the node in forwarding the packet. Specifically, in order to route a packet from node s to node t , each node in the graph demands some payment commensurate with the cost it incurs to handle the packet. The minimum cost path is chosen as the route, each node along the path getting the payment it demanded. Unfortunately, in most cases, the actual cost incurred is information private to the community owning the node and the protocol must assume that the community sets its own price. This can lead to cheating: communities will tend to inflate their operating costs to maximize the benefits received, leading to instability in the protocol. Thus, the protocol must be designed so that individual communities have no incentive to cheat. Such a *truthful mechanism* [1,3,4] will ensure that each community will demand a payment equal to its actual cost. The VCG mechanism [4,5,6,7] implements a truthful mechanism: the chosen route is the minimum cost according to the demanded payments, and each community gets paid the maximum amount it could have demanded to still be part of the chosen route, all other communities' demands remaining the same.

Since VCG is truthful, the chosen route is indeed the cheapest path with respect to the true cost. However, the payment made to the communities can be significantly greater than the solution cost. Hence, one has to analyze the amount by which the mechanism overpays, called the *frugality* of the mechanism [8,9,10]. This is measured by the *frugality ratio*, the maximum over all source-sink pairs of the ratio of the total payment made to the actual cost of the route.

The VCG mechanism and associated FR have been studied for shortest path routing on graphs, where each node or edge is considered an independent agent. We demonstrate in this work that the mechanism extends to the presence of *communities*. This captures the real-world nature of ad-hoc networks where nodes are organized into communities acting together, for example mobile users who group together following common social interests [11,12,13]. While this extension is simple for the standard VCG mechanism, we show that many natural extensions to VCG that remain computationally tractable in the usual case become intractable once communities are explicitly added to the model.

Random geometric graphs (RGG) [14] have been well-studied as theoretical models of ad-hoc networks [15,16,17,18]. Such graphs are constructed by placing nodes at random in the unit square, and adding an edge between two nodes if they are closer than the parameter r , which represents the broadcast radius. We consider various organizations of the nodes into k communities, including the traditional individual agent model in which each node is its own community (and $k = n$). We consider both the model where each node belongs to a uniformly at random selected community and the case where the node belongs to an arbitrary community (with no known underlying distribution). For any given community we assume that the per node cost is identical for all nodes of

the community. We take this to be a reasonable simplifying assumption which reflects the cooperative nature of nodes within a community, including that they may agree amongst themselves upon a fixed per node price. It may also reflect other forms of commonality of a given community's nodes, such as being of the same provider, being of the same general type, or sharing some locality in the clustered cases.

For a random geometric graph with k communities populated uniformly at random, where the costs are chosen uniformly at random from the interval $[c, c + B]$, we prove that the FR is bounded by $2\sqrt{2}(1 + \frac{2B(\log \log n)^2}{c \log n})$ w.h.p. For the individual node model (where each node is a different community), we show that the FR is bounded by $2(1 + \frac{B}{c})$ (respectively, $2(1 + \frac{B \log \log n}{c \log n})$) w.h.p. when costs are chosen arbitrarily (respectively, uniformly at random) from the interval $[c, c + B]$. Our proof techniques use the connectivity properties of RGG [15], together with iterated applications of the coupon collector's problem [19]. We also show a logarithmic bound in expectation when the number of communities in the network is small.

We also performed extensive network model simulations to see how VCG-based routing behaves in practice. The FR obtained in these simulations were always lower (better) than the theoretical upper bounds we provide. Our experiments also demonstrate that the FR goes up as the number of communities increase. This indicates that in the presence of many communities, a mechanism which minimizes the FR by weighting paths based on the number of communities may be desirable. In fact this is the intuition behind the result of [10] to improve over the FR of VCG. Unfortunately, we show that in the community model such weighting schemes become computationally intractable (NP-hard and even hard to approximate), implying that these improved mechanisms will be difficult to implement in practice. Due to space limitations, we have deferred the simulation results and some proofs to an extended version of the paper [20].

2 Related Work

The theory of algorithmic mechanism design was initiated by Nisan and Ronen in [4,21], in which they considered the generalized Vickrey-Clarke-Groves (VCG) mechanism [5,6,7] for various computational problems, including shortest path auctions. Although [4] considers VCG for general set systems, most subsequent work on truthful mechanisms for path auctions and the frugality thereof is restricted to the case where every edge is owned by an independent agent. Du et.al. [22] discuss a model where communities can own multiple edges, however in their model the identity of the community owning an edge is private, and they show that for such a model no truthful mechanism exists. In our work, we extend VCG for path auctions in the presence of communities where ownership is public but costs remain private. With the observation that VCG overpayments can be quite excessive for path auctions in worst cases, work has been put forth towards finding more frugal truthful mechanisms [8,9]. Karlin [10] proposed the \sqrt{n} mechanism, which is within a $\sqrt{2}$ factor of the frugality ratio for the best

truthful mechanism on any given graph, and in some cases performs up to $O(\sqrt{n})$ more frugally than VCG. In Section 6 we show that it is NP-hard to generalize many classes of truthful mechanisms for path auctions in the standard model, including the \sqrt{n} mechanism of [10], to the community model.

As we are interested in path auctions for ad-hoc networks, we study the performance of VCG for RGG [14], a model for the theoretical analysis of ad-hoc networks [15,16,17,18]. In particular, Gupta and Kumar [15] model ad-hoc networks as RGGs in their analysis of the critical radius required for asymptotic connectivity.

An alternative to the VCG is the first path auction where the agents on the winning path are paid their bid value. Immorlica et. al. [23] characterized all strong ϵ -Nash equilibria of a first path auction and showed that the total payment of this mechanism is often better than the VCG total payment. However, the drawback is that there is no guarantee that the bidders will reach an equilibrium, moreover, unlike the VCG, the preferred bid may depend on the communicating pair, which might not be known in advance.

VCG and variations thereof have been previously considered for routing in networks, fitting into a recent body of research tackling the problem of game-theoretic formalization of routing incentives for various networking domains [2,24,25,26]. Closest to our work in this regard is Andereg and Eidenbenz [2] paper in which they propose VCG for routing in ad-hoc networks. Although our work is nominally similar, there are crucial differences. In particular, while both consider VCG on ad-hoc networks, in their mechanism they consider nodes to have unbounded maximum potential radius, paying selected nodes to set their actual radius as desired according to how many bits they forward for the source-sink, and take each node to be an independent agent. We, on the other hand, consider a fixed topology in which radii are already set, and pay nodes to transmit according to some cost function set by their community.

Finally, we focus on previously unconsidered theoretical aspects of the problem, leaving the concrete implementation to a large body of work on implementation of internet currency [27] and other work dealing with the game-theoretic multi-hop routing [2,24,25,26] implementation.

3 Mechanism Design and the Payment Model

We model an ad-hoc network with k communities as a connected undirected graph $G = (V, E)$ where the nodes in V are partitioned into k subsets (the communities). Each community is assumed to be independently profit maximizing. We assume that there is no monopoly community in the graph, so that by removing one community from the graph the graph will still remain connected.

Given a k -community ad-hoc network (V, E) , and nodes $s, t \in V$, our goal is to design a protocol that will let s route a packet to t by a cheapest-cost path from s to t . A community i charges money for any packet that one of its node transmits. We assume all nodes in a community charge the same price, however, the exact determination of this cost is information private to the community.

While nodes can change location and connectivity over time, we assume that the network is static during the routing phase. We use tools from mechanism design [4] and define our protocol as follows.

1. We define a game on a k -community ad-hoc network (V, E) with k players, each corresponding to a community, and two states $s, t \in V$ (the source and the sink for routing). We define the *allowed outcomes* O of the game to be the finite set of simple paths between s and t .
2. For each path $o \in O$, each community i has a private cost $t^i(o)$ which is a function of the number of community nodes in path o and the cost of forwarding a packet by a node belonging to the community. We simplify the model by assuming that all the nodes belong to the same community have the same packet transmitting cost. Under this assumption $t^i(o) = C_i \cdot n_i(o)$, where C_i is the cost of transmitting one packet by a node of community i , and $n_i(o)$ is the number of i 's nodes lying on path o .
3. Each community defines a valuation function $p^i(o)$, which is the price it charges to transmit a packet on path o .
4. If the path \hat{o} is chosen as the route from s to t , then the utility function of community i will be $u^i(\hat{o}) = p^i(\hat{o}) - t^i(\hat{o})$ where $p^i(\hat{o}) \geq 0$ is the payment the community receives from the mechanism. The goal of community i is to maximize its utility $u^i(\hat{o})$.

The payment p^i to the communities is used to ensure a truthful implementation, i.e., an implementation where the dominant strategy of each community is to set its valuation p^i to be equal to t^i . We use the following payment in our mechanism. Let $d_{G|i=\infty}$ be the shortest path that does not contain any node belongs to community i and let $d_{G|i=0}$ be the cost of the shortest path where all nodes on the shortest path that belong to i have a zero cost. Then, the payment function $p^i(\hat{o}) = 0$ if i is not on the shortest path \hat{o} , and $p^i(\hat{o}) = d_{G|i=\infty} - d_{G|i=0}$ measures the maximum amount community i could have charged to still be part of the chosen route. This is a generalization of the shortest path payment scheme in [4]. Since shortest paths is a *monotone selection rule* (i.e., a losing community cannot become part of the shortest path by raising its valuation), standard techniques [4,8] show that this payment scheme implements a truthful mechanism. The *frugality ratio* is the ‘‘over payment’’ ratio of the mechanism:

$$FR = \frac{\sum_i p^i(\hat{o})}{\sum_i t^i(\hat{o})}.$$

4 Graph and Cost Model

A random geometric graph (RGG) with n nodes and radius r is constructed by picking n points (nodes) uniformly at random from the unit square, and putting an edge between nodes u and v if the distance between u and v is less than or equal to r .

Following previous theoretical work on ad-hoc networks [15], we represent ad-hoc networks as random geometric graphs. We choose the radius r at least on the order of asymptotic connectivity $r_{con} = \Omega(\sqrt{\frac{\log n}{n}})$ [15], i.e., the radius

that ensures that the graph is connected almost surely. Our models have four parameters: the number of nodes (n), the radius of the RGG (r), the number and choice of communities (k), and choice of transmission costs (F). We shall assume henceforth that $r \geq r_{con}$.

We consider three types of cost distribution functions F . First, we study *arbitrary bounded* cost distributions $F_A(c_{min}, B)$, where community picks an arbitrary cost from the interval $[c_{min}, c_{min} + B]$. As a special case, we study the *unit cost distribution* $F_C = F_A(1, 0)$ where each community charges unit cost per edge. Second, we study *uniformly-at-random bounded* cost distributions $F_U(c_{min}, B)$, where each community j picks a cost c_j uniformly at random from the interval $[c_{min}, c_{min} + B]$. Third, we study *uniformly-at-random unbounded* cost distributions $F_{A,U}(\epsilon)$, where $\epsilon > 0$, and each community j picks a cost c_j uniformly at random from the interval $[\epsilon, 1]$. As $\epsilon \rightarrow 0$, this model represents the case where the ratios of costs can be unbounded. Our worst case bounds depend on B , which becomes unbounded as $\epsilon \rightarrow 0$. While this is not a realistic case; it is interesting to see how bad the practical results can be. We study the following models:

Individual agent model. In the individual agent model (IAM), each node of the graph is its own community. This corresponds to the traditionally studied shortest path VCG mechanism on graphs where each node is an independent agent. We write $NC = (n, r, F)$ for an IAM network cost model with n nodes, radius r , and cost distribution F .

Random graph with communities. Given a number k of communities, each node in the random graph is assigned a community uniformly at random. We write $NC = (n, r, k, F)$ for the network cost model where there are n nodes, the radius is r , there are k communities (each node selecting its community uniformly at random), and the costs are determined according to the cost distribution F .

5 Theoretical Results

5.1 Frugality Ratio with High Probability

In many of the bounds, we use the following well known lemma on occupancy.

Lemma 1 (Balls in Bins [19,17]). *For a constant $c > 1$, if one throws $n \geq c\beta \log \beta$ balls into β bins, then w.h.p. both the minimum and the maximum number of balls in any bin is $\Theta(\frac{n}{\beta})$. Moreover, for $c < 1$ if one throws $n \leq c\beta \log \beta$ balls into β bins, then w.h.p. there will exist an empty bin.*

Due to the critical nature of the above threshold, we are able to give bounds w.h.p. for uniform distributions of costs and communities.

As mentioned previously, we consider random geometric graphs with radius chosen to guarantee connectivity w.h.p. Recall that we assume $r \geq r_{con}$. Although we shall state results for such general radii, we are primarily interested in small radii r such that $r = \Theta(r_{con})$. In particular, we will satisfy a slightly

stronger guarantee of *geo-denseness* [17], namely that, for any fixed arbitrary partitioning of the unit square into simple convex Euclidean regions β_i of area $\frac{r}{2\sqrt{2}} \times \frac{r}{2\sqrt{2}}$ each, every β_i will have the same order of nodes w.h.p. It follows from Lemma 1 that radius $\hat{r} = (2\sqrt{2} + \epsilon)\sqrt{\frac{\log n}{n}} \leq 3(r_{con,n})$ satisfies the geo-denseness property while still being on the same order as the radius for asymptotic connectivity. Henceforth, we will state some results for both general r and for \hat{r} as defined here. Note further that our following theoretical results hold for geo-dense geometric graphs in general, not only random geometric graphs. Due to space limitations, some proofs have been deferred to the full version of the paper.

Our first theorem considers the case of arbitrary costs in the Individual Agents Model (IAM), the standard model for path auctions.

Theorem 1 (IAM with Arbitrary Costs). *Given an IAM, $NC = (n, r, F_A(c_{\min}, B))$, for any $r \geq \hat{r}$, the FR of VCG is at most $2(1 + \frac{B}{c_{\min}})$ w.h.p.*

In particular, for IAM $NC = (n, r, F_C)$ with unit cost distribution, for any $r \geq \hat{r}$, the FR of VCG is at most 2. While unit costs do not seem to be a realistic assumption, and do not require notions of truthfulness, it yields insight into how the connectivity properties of a graph affect the overpayment. After all, with arbitrary costs one may obtain arbitrarily bad overpayments for any graph, but even with unit costs, the graph properties alone may yield bad overpayments. Therefore, the frugality ratio of VCG in the unit cost model is worthwhile to consider, and one that has been considered for other random graph models, namely Bernoulli graphs and random scale-free graphs, as well. A notable difference between random geometric graphs and those other two well-known random graph models is that while the hop diameter of the latter models is short w.h.p. the hop diameter of random geometric graphs is long w.h.p.

In standard shortest path auctions [4], unlike our model, costs are assigned on edges rather than nodes. For an IAM, $NC = (n, r, F_A(c_{\min}, B))$ where edge costs, we can similarly show that the FR is bounded by $2(1 + \frac{B}{c_{\min}})$ w.h.p.

When costs are distributed uniformly at random (u.a.r.), we may obtain provably better bounds than in the arbitrary case.

Theorem 2 (IAM with Random Costs). *Given $NC = (n, r, F_U(c_{\min}, B))$, for any $r \geq \hat{r}$, the FR is at most $2(1 + \frac{B}{bc_{\min}})$ where $b = \frac{\frac{nr^2}{8}}{2 \log(\frac{nr^2}{8})}$ w.h.p. In particular, for $r = \hat{r}$, if $B = O(c_{\min} \frac{\log n}{\log \log n})$, the FR of VCG for NC is a constant w.h.p.*

Now, we give our results for models with communities. The bounds of arbitrary costs are almost identical to that of the IAM.

Theorem 3 (Community Model with Arbitrary Costs). *Given $NC_C = (n, r, k, F_A(c_{\min}, B))$, for any $r \geq \hat{r}$, the FR is at most $2\sqrt{2}(1 + \frac{B}{c_{\min}})$ w.h.p.*

In particular, for $NC = (n, r, k, F_C)$, with unit costs, for any $r \geq \hat{r}$, the FR is at most $2\sqrt{2}$ w.h.p. Again, for the u.a.r. case, we obtain better guarantees.

Theorem 4 (Community Model with Random Costs). *Let $NC = (n, r, k, F_U(c_{\min}, B))$ with $r \geq \hat{r}$ and $k \leq \frac{8}{r^2}$ communities. For $b = \min\{\frac{k}{2 \log k}, \frac{nr^2}{2 \log \frac{nr^2}{8}}\}$ the FR of VCG is at most $2\sqrt{2}(1 + \frac{2B}{bc_{\min}})$ w.h.p. In particular, for $r = \hat{r}$ and $\log n \leq k \leq \frac{n}{\log n}$, if $B = O(c_{\min} \frac{\log n}{(\log \log n)^2})$, the FR is a constant w.h.p.*

Proof. Let s and t be an arbitrary source and sink pair and $SP = \langle v_0, v_1, \dots, v_d \rangle$ denote the shortest path between s and t . Since overpayments are made to communities rather than merely to nodes, partition SP into blocks $\langle L_1, \dots, L_q \rangle$ where each block belongs to a single community and consecutive blocks do not belong to the same community. For each community j , let $K_j = \langle L_{j_1}, \dots, L_{j_x} \rangle$ denote the set of blocks owned by community j . For each community j and block L_{j_i} denote by $v_{j_i,0}$ and $v_{j_i,f}$ the nodes in SP immediately preceding and succeeding L_{j_i} respectively, and let l_{j_i} be the line between $s' = v_{j_i,0}$ and $t' = v_{j_i,f}$. Partition l_{j_i} into $\frac{r}{2\sqrt{2}}$ length intervals (with at most one partial interval at the end of negligible effect) $y \in \{1, 2, \dots, \frac{d(s',t')}{\frac{r}{2\sqrt{2}}}\}$. Depending on how close l_{j_i} is to a boundary of the unit square, it is clear that there must exist a $\frac{r}{2\sqrt{2}} \times d(s', t')$ rectangular area A_{j_i} with l_{j_i} as one of the sides lying entirely inside the unit square. Depending on the orientation of this rectangular area, for each interval y , let S_y denote the $\frac{r}{2\sqrt{2}} \times \frac{r}{2\sqrt{2}}$ square in A_{j_i} with interval y as one of the sides.

By Lemma 1 and the choice of r , there are $\Theta(\frac{nr^2}{8})$ nodes in each S_y w.h.p. Each node chooses amongst the k communities u.a.r. Each of k communities chooses its cost u.a.r. from $[c_{\min}, \dots, c_{\min} + B]$. By the choice of b , w.h.p. the number of communities in each cost interval of the form $[c_{\min} + (\alpha - 1)\frac{B}{b}, c_{\min} + \alpha\frac{B}{b}]$ (for α from 1 to b) is $\Theta(\frac{k}{b})$. Therefore, since the number of communities in each cost interval is on the same order, each node in S_y picks amongst the cost intervals as well up to constant factors. Again, by the choice of b , the number of cost intervals and re-application of Lemma 1, we have that for each cost interval α there are $\Theta(\frac{nr^2}{8b})$ nodes of S_y having cost in interval α . Then, recalling that consecutive bins form a clique, we may route along nodes in the first two cost intervals in each square bin, depending upon which cost interval the corresponding community in SP lies. Then, for each A_{j_i} , we obtain a path of cost at most $2\sqrt{2}\frac{d(v_{j_i,0}, v_{j_i,f})}{r}(c_{\min} + \frac{2B}{b})$ other than L_{j_i} which has cost at least $d(v_{j_i,0}, v_{j_i,f})rc_{\min}$. So, for L_{j_i} , the FR is at most $2\sqrt{2}\frac{2B+c_{\min}}{bc_{\min}}$. Summing over each L_{j_i} , we obtain the same ratio. This characterizes the payment to community j . Moreover, the argument is the same for any community since the scaling by distance is lost. Thus, the theorem follows. \square

5.2 Frugality Ratio in Expectation

The bounds so far are all with high probability. However, in the case of fewer communities we may find significantly improved bounds of VCG with communities for RGGs *in expectation*. When the number of communities k is $O(\frac{\log n}{\log \log n})$

(or, for general r , when k is $O(\frac{nr^2}{\log(nr^2)})$) we may note once again that every community occurs in every bin (of $\frac{r}{2\sqrt{2}}x\frac{r}{2\sqrt{2}}$ size). So, due to the aforementioned bin properties for RGGs, we need only bound the expected ratio of the second cheapest community to the cheapest community.

Theorem 5. *Let $NC = (n, r, k, F_U(c_{\min}, B))$ with radius $r \geq \hat{r}$ and $k \leq \frac{nr^2}{\log(nr^2)}$ communities. The expected FR of VCG for NC is $O(\min\{\log \frac{B}{c_{\min}}, \frac{B}{kc_{\min}}\})$ w.h.p.*

Proof. Due to aforementioned geometric bin properties and normalization, it suffices to show that the expected ratio of the second cheapest to the cheapest of k costs chosen u.a.r from $[1, B]$ is $O(\log B)$. As such, note that the probability that the cheapest is in $[x, x + dx]$ is $k\frac{dx}{B-1}(\frac{B-x}{B-1})^{k-1}$, corresponding to the choices for the cheapest variable and the event that that variable is in $[x, x + dx]$ and all rest are in $(x, B]$. Moreover, the expected value of the second cheapest given that the cheapest is x is the expected value of the cheapest of the $k-1$ restricted to interval $(x, B]$, which is easy to check to be $\frac{B+x(k-1)}{xk}$. Thus,

$$\begin{aligned} E_k[\frac{Y}{X}] &= \int_1^B \frac{k}{B-1} (\frac{B-x}{B-1})^{k-1} \frac{1}{x} \frac{B+x(k-1)}{xk} dx = \frac{B}{B-1} ((\int_1^B (\frac{B-x}{B-1})^{k-1} \frac{dx}{x}) + \frac{k-1}{k}) \\ &\leq \frac{B}{B-1} (\min\{\log B, \frac{B-1}{k}\} + \frac{k-1}{k}) \end{aligned}$$

□

We may generalize the expected ratio of the second cheapest to the cheapest of k i.i.d. random costs given cumulative distribution F and density function f as follows: The probability that the minimum is in $[x, x + dx]$ is, taking over the k choices of the minimum variable, $kf(x)(1-F(x))^{k-1}$. Similarly, the probability that the second cheapest is in $[y, y + dy]$ given that the cheapest is x is the probability that the minimum of the remaining $k-1$ is in $[y, y + dy]$ given that all $k-1$ have cost greater than x . Thus, the expectation in question is:

$$\begin{aligned} E_k[\frac{Y}{X}] &= \int_1^\infty \frac{kf(x)(1-F(x))^{k-1}}{x} dx \int_x^\infty y(k-1) \frac{f(y)(1-F(y))^{k-2}}{(1-F(x))^{k-1}} dy \\ &= k(k-1) \int_1^\infty \frac{f(x)}{x} dx \int_x^\infty yf(y)(1-F(y))^{k-2} dy \end{aligned}$$

Substituting, we obtain the following results for some

Corollary 1. *Let $NC_\lambda = (n, r, k, F_\lambda, B)$ with $r \geq \hat{r}$ and $k \leq \frac{nr^2}{\log(nr^2)}$ communities and F_λ the exponential distribution translated by $+1$ with parameter λ . The expected FR of VCG for NC_λ is at most $4\sqrt{2}$ w.h.p..*

For the distribution F_{recip} obtained by taking reciprocals of random variables chosen according to the uniform distribution on the unit interval $(0, 1]$, in the model $NC_{recip} = (n, r, k, F_{recip}, B)$ with radius $r \geq \hat{r}$ and $k \leq \frac{nr^2}{\log(nr^2)}$ communities, we similarly get that the expected FR of VCG for NC_{recip} is $2\sqrt{2}\frac{k-1}{k-2}$ w.h.p. In fact, we can say something much stronger for this distribution.

Lemma 2. *Let $NC_{recip} = (n, r, k, F_{recip}, B)$ with radius $r \geq \hat{r}$ and $k \geq nr^2$ communities. The FR of VCG for NC_{recip} is at most $2e^3\sqrt{2}$ w.h.p..*

This holds because, by the geometric bin properties, it suffices to show that *within each bin* the probability that the second cheapest in that bin is more than e^3 times the cheapest in that bin is $O(\frac{1}{nm})$, where $m = \frac{8}{r^2}$ is the number of bins. Let q denote the number of communities occurring w.h.p. in every bin. By choice of k and r , we have $q = \Theta(\frac{nr^2}{8})$ by coupon collection. The event that $\frac{1}{X} \geq e^3 \frac{1}{Y}$ implies that $q - 1$ reciprocals chosen u.a.r. all lay in $(0, \frac{1}{e^3})$, the probability of which is $\frac{q}{e^{3(q-1)}}$. Thus, $Pr[\frac{Y}{X} \geq e^3] = Pr[\frac{1}{X} \geq e^3 \frac{1}{Y}] < \frac{q}{e^{3(q-1)}}$, where X is the cheapest and Y is the second cheapest. Moreover, $\frac{q}{e^{3(q-1)}} \leq \frac{q}{n^2} = \frac{r^2}{n}$ by choice of q , completing the proof.

By noting that, for F_λ , the exponential distribution translated by $+1$, the probability that $q - 1$ costs are higher than A is at most $ke^{-\lambda(A-1)(q-1)}$, a very similar argument gives the following.

Lemma 3. *Let $NC_\lambda = (n, r, k, F_\lambda, B)$ with radius $r \geq \hat{r}$ and $k \geq nr^2$ communities. The FR of VCG for NC_λ is $O(1)$ w.h.p..*

6 Hardness of Extensions

NP-Hardness of Extensions. Both simulation results and related work on the traditional path auction model [8,9,28,10] suggest that a mechanism that minimizes some weighting of total path costs by the number of communities on the path may have a lower FR than VCG. For example, the mechanism proposed in [10] is known to be up to \sqrt{n} times more frugal than VCG. Unfortunately, as we show next, in the presence of communities, the implementation of this mechanism requires solving intractable problems.

The first step of the \sqrt{n} mechanism of [10] is to find the least cost edge-disjoint cycle through s and t . In the community model, this would correspond to finding at least some community disjoint cycle through s and t . Note that the existence of two community disjoint paths is not guaranteed by the no-monopoly condition. For example, consider $k = 3$ and a graph consisting of three length paths P_1, P_2, P_3 from s to t where each path P_i excludes only community i .

By representing each community with a unique color, we color the nodes (or, alternately, edges, the results apply to both cases) according to their communities. Finding a community disjoint cycle is the same as finding a color-disjoint cycle. This problem is NP-Complete by a reduction from 3-SAT. A similar problem is independently shown to be NP-Complete in [29].

Lemma 4. *Consider the problem \mathfrak{C} : Given a graph $G = (V, E)$ with nodes arbitrarily colored from k colors, and a designated source-sink pair (s, t) , find two color disjoint paths through s and t . \mathfrak{C} is NP-Complete. The same is true considering edge colorings instead of node colorings.*

APX-Hardness of Natural Extensions. As a second possible extension, we can study the VCG under other cost models. For example, we could try to minimize the number of communities along the shortest path in order to try to

reduce the FR. Unfortunately, we have found that many approaches in these directions turn out to be NP-complete, some even strongly approximation hard.

Here we show that any natural truthful mechanism with a selection rule incorporating some kind of minimization of the number of communities on the path is strongly approximation-hard to compute. Our reduction is an approximation preserving reduction from the Minimum Monotone Satisfying Assignment ($MMSA_3$) problem, which is known to be $2^{\log^{1-o(1)} n}$ hard to approximate [30,31]. While there are closely related approximation hardness results under various names [32,33], our result and reduction are both more general and more direct. First, notice that for all $0 < x < 1$, we have $2^{\log^{1-o(1)} n} > n^x$. Now, we define a natural class of truthful mechanisms for path auctions in the community model in the following way. A truthful mechanism for path auctions in the community model (with per unit costs) is a (f, g) *min-agent mechanism* if its monotonic selection rule is of the following form. Given source s and destination t , select the path P from s to t that minimizes the product $f(q)g(p)$, for some strictly increasing, efficiently invertible function f and non-decreasing function g , where q is the number of communities on P and p is the total cost of P . Now, we proceed to our hardness result.

Theorem 6. *For any $0 < x < 1$, for any increasing, efficiently invertible function f and non-decreasing function g , the selection rule of a (f, g) min-agent mechanism is $f(k_n^x)$ hard to approximate, where k_n is the total number of communities and n is the number of nodes.*

The same proof also implies the approximation-hardness of even computing VCG for various other cost-functions involving the community model, such as fixed community-network entrance fees (i.e., a one-time fee C_i for using any number of community i 's nodes, which may be a more natural model for some service providers). The following is obtained by taking g to be a constant function in Theorem 6.

Corollary 2. *VCG for the community model under fixed community sub-network entrance fees is hard to approximate to within k^x , for any $0 < x < 1$, given k total communities.*

Acknowledgments. We thank Deborah Estrin for many helpful discussions about models for ad-hoc networks. We also thank Chen Avin, Ilya Shpitser, and Sandra Batista for helpful discussions and comments. This research was sponsored in part by the NSF grants CCF-0427202 and CCF-0546170.

References

1. Papadimitriou, C.: Algorithms, games, and the internet. In: STOC (2001)
2. Anderegg, L., Eidenbenz, S.: Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents. In: MOBICOM (2003)
3. Kreps, D.M.: A Course in Microeconomic Theory. Princeton University Press, Princeton (1990)
4. Nisan, N., Ronen, A.: Algorithmic mechanism design. Games and Economic Behavior 35, 166–196 (2001)

5. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16, 8–37 (1961)
6. Clarke, E.: Multipart pricing of public goods. *Public Choice* 11, 17–33 (1971)
7. Groves, T.: Incentives in teams. *Econometrica* 41, 617–631 (1973)
8. Archer, A., Tardos, E.: Frugal path mechanisms. In: *SODA* (2002)
9. Talwar, K.: The price of truth: Frugality in truthful mechanisms. In: *STACS* (2003)
10. Karlin, A., Kempe, D., Tamir, T.: Beyond VCG: Frugality of truthful mechanisms. In: *FOCS* (2005)
11. Kortuem, G., Segall, Z.: Wearable communities: Augmenting social networks with wearable computers. *IEEE Pervasive Computing Magazine* 2(1), 71–78 (2003)
12. Schulz, S., Herrmann, K., Kalckloesch, R., Schwotzer, T.: Towards trust-based knowledge management in mobile communities. In: *IAAA* (2003)
13. Musolesi, M., Hailes, S., Mascolo, C.: An ad hoc mobility model founded on social network theory. In: *MSWiM* (2004)
14. Penrose, M.D.: *Random Geometric Graphs*. Oxford University Press, Oxford (2003)
15. Gupta, P., Kumar, P.: The capacity of wireless networks. *IEEE TIT*, Los Alamitos (2000)
16. Goel, A., Rai, S., Krishnamachari, B.: Monotone properties of random geometric graphs have sharp thresholds. *Annals of Applied Probability* 15 (2005)
17. Avin, C., Ercal, G.: On the Cover Time of Random Geometric Graphs. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) *ICALP 2005*. LNCS, vol. 3580, pp. 677–689. Springer, Heidelberg (2005)
18. Diaz, J., Petit, J., Serna, M.J.: Faulty random geometric networks. *PPL* (2000)
19. Motwani, R., Raghavan, P.: *Randomized algorithms*, Cambridge (1995)
20. Ercal, G., Izhak-Ratzin, R., Majumdar, R., Meyerson, A.: Frugal routing on wireless ad-hoc networks. Technical report, University of California, Los Angeles (2008)
21. Nisan, N., Ronen, A.: Computationally feasible VCG mechanisms. In: *EC* (2000)
22. Du, Y., Sami, R., Shi, Y.: Path auction games when an agent can own multiple edges. In: *NetEcon* (2006)
23. Immorlica, N., Karger, D., Nikolova, E., Sami, R.: First-price path auctions. In: *EC* (2005)
24. Sun, H., Song, J.: Strategy proof trust management in wireless ad hoc network. In: *Electrical and Computer Eng'* (2004)
25. Wang, W., Li, X.-Y., Eidenbenz, S., Wang, Y.: Ours: optimal unicast routing systems in non-cooperative wireless networks. In: *MobiCom* (2006)
26. Zhong, S., Li, L.E., Liu, Y.G., Yang, Y(R.): On designing incentive-compatible routing and forwarding protocols in wireless ad-hoc networks: an integrated approach using game theoretical and cryptographic techniques. In: *MobiCom* (2005)
27. Buttyán, L., Hubaux, J.: A virtual currency to stimulate cooperation in self-organized ad hoc networks. Technical Report DSC (2001)
28. Elkind, E., Sahai, A., Steiglitz, K.: Frugality in path auctions. In: *SODA* (2004)
29. Yuan, S., Jue, J.P.: Dynamic lightpath protection in WDM mesh networks under wavelength-continuity and risk-disjoint constraints. *Comput. Netw* (2005)
30. Dinur, I., Safra, S.: On the hardness of approximating label-cover. In: *IPL* (2004)
31. Alekhovich, M., Buss, S.R., Moran, S., Pitassi, T.: Minimum propositional proof length is NP-hard to linearly approximate. In: *MFCS* (1998)
32. Carr, R.D., Doddi, S., Konjevod, G., Marathe, M.: On the red-blue set cover problem. In: *SODA* (2000)
33. Wirth, H.C.: *Multicriteria Approximation of Network Design and Network Upgrade Problems*. PhD thesis (2001)