Hindawi Publishing Corporation
International Journal of Distributed Sensor Networks
Volume 2011, Article ID 895398, 17 pages
doi:10.1155/2011/895398

Research Article

VCG with Communities on Random Ad Hoc Networks

Gunes Ercal,[1] Rafit Izhak-Ratzin,[2] Rupak Majumdar,[3] and Adam Meyerson[4]

[1] *University of Kansas, Lawrence, KS 66045, USA*
[2] *Palo Alto Networks, Sunnyvale, CA 94089, USA*
[3] *Max Planck Institute, 66123 Saarbruecken, Germany*
[4] *University of California, Los Angeles, CA 90095, USA*

Correspondence should be addressed to Gunes Ercal, gunes@ku.edu

Received 17 October 2010; Revised 24 February 2011; Accepted 7 April 2011

We study game-theoretic mechanisms for routing in wireless ad hoc networks. Our major results include a combination of theoretical bounds and extensive simulations, showing that VCG-based routing in wireless ad-hoc networks exhibits small frugality ratio with high probability. Game-theoretic mechanisms capture the noncooperative and selfish behavior of nodes in a resource-constrained environment. There have been some recent proposals to use these 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 are the first to show frugality bounds for random geometric graphs, a well-known model for ad-hoc wireless connectivity. In addition, we generalize the model of agent behavior by allowing sets of nodes to form *communities* to maximize total profit. We are the first to analyze the performance of VCG under such a community model. 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

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; a node will always agree to forward a packet to its recipient. In resource-scarce environments (prevalent in ad-hoc networks), this co-operativeness assumption is suspect. Forwarding a packet incurs some cost to the node (i.e., the use of battery power), 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 often 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, and 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. This simplifies the protocol design by eliminating the need to model each community's knowledge of each other or their cost distributions. For network routing, the VCG mechanism [4–7] 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, there is a cost to truthfulness. The payment made to the communities can be significantly greater than the cost of the solution. Hence, one has to analyze the the amount by which the mechanism overpays, called the *frugality* of the mechanism [8–10]. This is measured by the *frugality ratio*, the maximum over all source-sink pairs of the ratio of the total payment made to the cost of the route.

The VCG mechanism and associated frugality ratio have been studied for shortest path routing on graphs, where each node or edge is considered an independent agent. In particular, some different routing protocols for wireless ad-hoc networks that are based on variations of the VCG mechanism have been proposed by the research community [2, 11–13]. We demonstrate in this work that the mechanism extends to the presence of *communities* (i.e., where nodes of the graph are partitioned into independent profit-making agents). This captures the real-world nature of ad-hoc networks where nodes are organized into companies or communities acting together, for example, mobile users who group together following common social interests [14–17], mobile users that have a relation of trust [18], and so forth. 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.

We study the frugality ratio (a measure of cost-efficiency) of the generalized VCG mechanism for reliable routing in the presence of noncooperative behavior on wireless ad-hoc networks. We provide theoretical bounds on the frugality ratio and additionally validate VCG-based routing with extensive simulations. We are the first to demonstrate frugal bounds on VCG both for random geometric graphs and under the generalized community model.

Random geometric graphs [19] 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. Such graphs have been well studied as theoretical models of wireless ad-hoc networks [20–25]. 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 $k = n$). In the theoretical bounds, 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). In simulations, we additionally consider more clustered situations where the communities have some locality. While we consider both the cases where per-node costs amongst different communities are distributed uniformly at random and distributed arbitrarily, 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 frugality ratio is bounded by $2\sqrt{2}(1 + 2B(\log\log n)^2/c\log n)$ with high probability. For the (traditionally studied) individual node model (where each node is a different community), we show that the frugality ratio is bounded by $2(1 + B/c)$ (resp., $2(1 + B\log\log n/c\log n)$) with high probability when costs are chosen arbitrarily (resp., uniformly at random) from the interval $[c, c + B]$. Our proof techniques use the connectivity properties of random geometric graphs [20], together with iterated applications of the coupon collector's problem [26]. We also show a logarithmic bound in expectation when the number of communities in the network is small.

We also perform a detailed study of VCG-based routing and its frugality ratio through extensive simulation of network models. We simulate random geometric graphs with few ($O(\log n/\log\log n)$) big communities as well as many ($O(n/\log n)$) small communities. Our simulations show that the frugality ratio obtained is lower (better) than the theoretical upper bounds we provide. In addition, we simulate networks where our theoretical results will not directly apply; these include networks consisting of some combination of large and small communities, and networks where nodes belonging to a particular community exhibit geographic locality. In all cases, the frugality ratios remain low. We note that such extensive simulations properly extend our conference paper [27] which included only theoretical bounds.

Our experiments demonstrate that the frugality ratio goes up as the number of communities increases. This indicates that in the presence of many communities, a mechanism which explicitly minimizes the frugality ratio 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 frugality ratio 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.

While we only concentrate on the frugality of VCG mechanisms in this paper, our protocols can be implemented for real-world networks using the techniques of [2], using real or virtual payments [28].

2. Related Work

The theory of algorithmic mechanism design was initiated by Nisan and Ronen in [4, 29], in which they considered the generalized Vickrey-Clarke-Groves (VCG) mechanism [5–7] for various computational problems, including shortest path auctions. Since then, 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]. Much of this research has resulted in similarly worst case bounds for any truthful mechanism [30]. However, the appropriate question to ask is whether we can define a mechanism *for a particular graph* which comes close to the best possible mechanism *for that graph*. To resolve this question, [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. In some cases, this performs up to $O(\sqrt{n})$ more frugally than VCG. Although Nisan and Ronen's original paper 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. [31] 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. In Section 7, we show that it is NP-hard to generalize important classes of truthful mechanisms for path auctions in the standard model, including the $\sqrt{n}$ mechanism, to the community model.

Although VCG performs badly in worst cases, it has been observed that VCG yields small frugality ratio for various random graph models such as random Bernoulli graphs and random scale-free graphs [32–34]. Here, as we are interested in path auctions for ad-hoc and wireless networks, we study the performance of VCG for random geometric graphs [19], a classical model for wireless networks that has been considered in much theoretical work in this area [20–25]. In particular, it is the model that has been considered by Gupta and Kumar [20] in their analysis of the critical radius required for asymptotic connectivity in such networks, and in this work, we consider radii on this order as we require connectivity and small radius. Random geometric graphs have repeatedly been found to share similar threshold properties with random Bernoulli graphs [19, 21], such as asymptotic connectivity probability and radius [20], optimal cover time [23], and, as we show in this work, constant frugality ratio for VCG under unit costs in the standard model. However, despite the similarities in resulting thresholds, proof and analysis of these thresholds have required strikingly different methodologies due to inherent sharp differences between random Bernoulli graphs and random geometric graphs [21, 23]. Similarly, here we may not utilize previous methods of bounding VCG overpayments derived from the results on random Bernoulli graphs (or random scale-free graphs) as that analysis relies heavily on expansion properties or short diameter of such graphs, neither of which is shared by random geometric graphs.

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. in [35] 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 of this mechanism 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 reliable routing incentives for various networking domains, such as peer-to-peer networks and ad-hoc networks [2, 11–13], only [2]. Closest to our work in this regard is the paper of Anderegg and Eidenbenz [2] in which they propose VCG for routing in ad-hoc networks. Although our work is nominally similar, there are crucial differences. In particular, although both works 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 (one may view this as assuming bounded transmission radius) and pay nodes to transmit according to some cost function set by the community that the nodes belong to taking into account various factors (e.g., energy, quality of service, etc.). Our work fits a more general framework for routing mechanisms where we do not assume to know exactly how forwarding costs are dictated. Also, our analysis of frugality ratios considers both arbitrary and random cost distributions both in the presence of communities and for individual independent agents, and in simulations, we further consider the case of clustered graphs as well.

Finally, we focus on previously unconsidered important theoretical aspects of the problem in this work, leaving the concrete implementation to a large body of work on implementation of internet currency [28] and other previous work dealing extensively with the implementation of game-theoretic multihop routing [2, 11–13]. We reiterate both our results on NP-hardness and APX-hardness of natural extensions as well as previous work on impossibility of ensuring truthfulness for some extremely generalized agent models [31]. In light of these, as well as our low frugality bounds for VCG obtained via proofs and simulations, we recommend VCG as both a reliable and *cost-efficient* routing protocol for wireless ad-hoc networks under reasonable generalized (in comparison to traditional models) modes of node behavior, selfishness, and cooperation.

3. The Payment Model

In this section, we describe our 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 and t from 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.

Costs are incurred by the nodes in forwarding packets. A community i charges money for any packet that one of its node transmits. This cost reflects, for example, the power and other resources required to forward a packet, as well as other factors like the location and number of agents belong to the community. We assume all nodes of 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.

One protocol is to let each community declare its true cost and then find a shortest path in the resulting graph, each community along the path getting paid the amount it had demanded for each of its nodes in the path. However, since communities want to maximize their profit independently, they might inflate their actual cost in order to maximize their payment. Hence, what we want is a protocol that provides no incentive for cheating. We use tools from mechanism design [3, 4, 29] to design such a *truthful* mechanism.

We define our protocol as a mechanism design problem as follows.

(1) We define a game on a k-community ad-hoc network (V, E) with k players, each corresponding to a community. 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. This is private information for the community: all and only the nodes in community i know the function, $t^i(o \in O)$. 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 (private) valuation function $t^i(o)$, which is the price it charges to transmit a packet on path o.

(4) If the path o is chosen as the route from s to t, then the utility function of community i will be $u^i(o) = p^i(o) - t^i(o)$ where $p^i(o) \geq 0$ is the payment the community receives from the mechanism. The goal of community i is to maximize its utility $u^i(o)$.

The payment p^i to the communities is used to ensure a truthful implementation, that is, an implementation where the dominant strategy of each community is to set its valuation v^i to be equal to t^i. Such a truthful mechanism is the Vickrey-Groves-Clarke (VCG) mechanism [5–7]. We use the following payment (which is an easy generalization of the payment scheme for shortest paths on graphs studied in [4]) in our mechanism. Let $d_{G|i=\infty}$ be the shortest path that does not contain any node belonging 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 $p^i(o) = 0$ if i is not on the shortest path o, and $p^i(o) = d_{G|i=\infty} - d_{G|i=0}$ measures the maximum amount community i could have charged to still be part of the chosen route, namely, the *threshold bid* for that community. Since shortest path 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, that is, there is no incentive for any community to lie about its cost.

The *frugality ratio* is the "overpayment" ratio of the mechanism. Since VCG selects the shortest path o, the frugality ratio will be $\mathrm{FR} = \sum_i p^i(o)/\sum_i t^i(o)$.

4. Graph and Cost Model

We represent ad-hoc and wireless networks as random geometric graphs with radius at least on the order of asymptotic connectivity $r_{\mathrm{con}} = \Omega(\sqrt{\log n/n})$ [20]. To generate a random geometric graph with n nodes and radius r, n points (nodes) are picked uniformly at random from the unit square, and there is an edge between nodes u and v if the distance between u and v is less than or equal to r. As the random geometric graph is a standard model in theoretical work on ad-hoc and wireless networks [20], and as connectivity is a minimum requirement for any routing in a network, we take these to be reasonable assumptions.

Our models have four parameters: the number of nodes (n), the radius of the random geometric graph (r), the number and choice of communities (k), and choice of transmission costs (F).

As explained above, we use $r \geq r_{\mathrm{con}}$. In simulations, we actually use $r = \Theta(r_{\mathrm{con}}(n))$ to study the behavior of the network with the minimal radius that ensures network connectivity. Our theoretical results are for general $r \geq r_{\mathrm{con}}$.

We consider 3 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 of unbounded differences in costs. Our worst-case bounds depend on B, which becomes unbounded as $\epsilon \rightarrow 0$. While this is probably not a realistic case, we find it interesting to see how bad the practical results can be.

In our theoretical results and our simulations, we study the following models (obtained by varying the parameters) in our paper. The models and the results are summarized in Table 1.

Individual Agent Model. In the individual agent model (IAM), each node of the graph is its own community. This

Table 1: Summary of all the models that we consider in this paper. "uar bdd" (resp., "uar unbdd") is the uniformly at random bounded (resp., uniformly at random unbounded) cost distribution.

Community Model	Random G uar bdd	Random G uar unbdd	Clustured G uar bdd
Individual Agent Model	Section 6.1, Theorem 5.3	Section 6.2	Section 6.3
Small number of large Comp', $k = \log n/\log\log n$	Section 6.1, Theorems 5.5 and 5.6	Section 6.2, Corollary 5.8	Section 6.3
Mixed model	Section 6.1	Section 6.2	Section 6.3
Large number of small Comp', $k = n/\log n$	Section 6.1, Theorem 5.5	Section 6.2, Lemma 5.9	Section 6.3

corresponds to the traditionally studied shortest path VCG mechanism on graphs where each node is an independent agent. We provide theoretical bounds on the frugality ratio for the IAM for random geometric graphs with both arbitrary bounded cost distributions and uniformly-at-random bounded cost distributions. We write NC $= \{n, r, F\}$ for an individual agent 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. We provide theoretical upper bounds on the frugality ratio for $\{n, r, k, F_A(c_{\min}, B)\}$ where $r \geq r_{\text{con}}$, as well as for $\{n, r, k, F_U(c_{\min}, B)\}$ where $r \geq r_{\text{con}}$ and $k \leq 8/r^2$.

Further, for our simulation, we study three different cases: a small number of large communities corresponding to $k = \Theta(\log n/\log\log n)$, a large number of small communities corresponding to $k = \Theta(n/\log n)$, and finally, a mixed model with a small number of large communities and larger number of small communities. In the last model, a node will choose with high probability to be in one of the large communities and with small probability to be in one of the small communities.

Clustered Graph. In order to evaluate the VCG mechanism in the presence of real-world structures, we also simulate a clustered model that reflects geographical structure in the real world. In the clustered model, each community i chooses a center cnt_i uniformly at random. There is a fixed radius r_t. Each node v belongs to community i is chosen uniformly at random from within the circle centered at cnt_i with radius r_t. This represents geographical locality common in real networks.

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 5.1 (balls in bins [23, 26]). *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 $\Theta(n/\beta)$. 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 with high probability for uniform distributions of costs and communities.

As mentioned previously, we consider random geometric graphs with radius chosen to guarantee connectivity with high probability. Recalling that $r_{\text{con},n} = \sqrt{(\log n + \gamma_n)/n}$, for any increasing function γ_n, is the critical radius for asymptotic connectivity [20], we require that $r \geq r_{\text{con},n}$. Although we will state results for such general radii, we are primarily interested in small radii r such that $r = \Theta(r_{\text{con},n})$. In particular, we will satisfy a slightly stronger guarantee of *geo-denseness* [23], namely that, for any fixed arbitrary partitioning of the unit square into simple convex Euclidean regions β_i of area $(r/2\sqrt{2}) \times (r/2\sqrt{2})$ each, every β_i will have the same order of nodes with high probability. It follows from Lemma 5.1 that radius $\hat{r} = (2\sqrt{2} + \epsilon)\sqrt{\log n/n} \leq 3(r_{\text{con},n})$ satisfies the geodenseness 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 geodense geometric graphs in general, not only random geometric graphs. In order to maintain continuity, most proofs have been deferred to the appendix section.

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

Theorem 5.2 (IAM with arbitrary costs). *Given an IAM NC $= \{n, r, F_A(c_{\min}, B)\}$, for any $r \geq \hat{r}$, the frugality ratio of VCG is at most $2(1 + B/c_{\min})$ w.h.p.*

In particular, for IAM NC $= \{n, r, F_C\}$ with unit cost distribution, for any $r \geq \hat{r}$, the frugality ratio 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 the costs are on edges, we can similarly show that the frugality ratio is bounded by $2(1 + B/c_{\min})$ w.h.p.

When costs are distributed uniformly at random (i.e., under the cost model $F_U(c_{\min}, B)$), we may obtain provably better bounds than in the arbitrary case.

Theorem 5.3 (IAM with random costs). *Given NC* $= \{n, r, F_U(c_{\min}, B)\}$, *for any* $r \geq \hat{r}$, *the frugality ratio is at most* $2(1 + B/bc_{\min})$ *where* $b = (nr^2/8)/2\log(nr^2/8)$ *w.h.p. In particular, for* $r = \hat{r}$, *if* $B = O(c_{\min}\log n/\text{loglogn})$, *the frugality ratio 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 5.4 (community model with arbitrary costs). *Given* $NC_C = \{n, r, k, F_A(c_{\min}, B)\}$, *for any* $r \geq \hat{r}$, *the frugality ratio is at most* $2\sqrt{2}(1 + B/c_{\min})$ *w.h.p.*

In particular, for NC $= \{n, r, k, F_C\}$, with unit costs, for any $r \geq \hat{r}$, the frugality ratio is at most $2\sqrt{2}$ w.h.p. Again, for costs distributed uniformly at random, we obtain better guarantees.

Theorem 5.5 (community model with random costs). *Let NC* $= \{n, r, k, F_U(c_{\min}, B)\}$ *with radius* $r \geq \hat{r}$ *and* $k \leq 8/r^2$ *communities. For*

$$b = \min\left\{\frac{k}{2\log k}, \frac{nr^2/8}{2\log(nr^2/8)}\right\}, \tag{1}$$

the frugality ratio of VCG is at most $2\sqrt{2}(1 + 2B/bc_{\min})$ *w.h.p. In particular, for* $r = \hat{r}$ *and* $\log n \leq k \leq n/\log n$, *if* $B = O(c_{\min}(\log n/(\log\log n)^2))$, *the frugality ratio is a constant w.h.p.*

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

By Lemma 5.1 and the choice of r, there are $\Theta(nr^2/8)$ nodes in each S_y w.h.p. Each node chooses amongst the k communities uniformly at random. Each of k communities chooses its cost uniformly at random from $[c_{\min}, \ldots, c_{\min} + B]$. By the choice of b, w.h.p. the number of communities in each cost interval of the form $[c_{\min} + (\alpha - 1)(B/b), c_{\min} + \alpha(B/b)]$ (for α from 1 to b) is $\Theta(k/b)$. 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 uniformly at random as well up to constant factors. Again, also by the choice of b, the number of cost intervals and reapplication of Lemma 5.1, we have that for each cost interval α there are $\Theta(nr^2/8b)$ 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 $2\sqrt{2}(d(v_{j_i,0}, v_{j_i,f})/r)(c_{\min}+2B/b)$ 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 frugality ratio is at most $2\sqrt{2}((2B + c_{\min})/bc_{\min})$. Summing over each L_{j_i}, we just obtain the same ratio. This characterizes the payment to community j. Moreover, clearly, the argument is the same for any community since the scaling by distance is lost. Thus, the theorem follows. □

5.2. Frugality Ratio in Expectation. The bounds so far are all with very 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 $O(\log n/\log\log n)$ (or, for general r, when k is $O(nr^2/\log(nr^2))$), we may note once again that every community occurs in every bin (of $(r/2\sqrt{2}) \times (r/2\sqrt{2})$ size). So, due to the aforementioned bin properties for RGGs, we need only to bound the expected ratio of the second cheapest community to the cheapest community.

Theorem 5.6. *Let NC* $= \{n, r, k, F_U(c_{\min}, B)\}$ *with radius* $r \geq \hat{r}$ *and* $k \leq nr^2/\log(nr^2)$ *communities, then the expected frugality ratio of VCG for NC is* $O(\min\{\log(B/c_{\min}), B/kc_{\min}\})$ *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 uniformly at random from $[1, B]$ is $O(\log B)$. As such, note that the probability that the cheapest is in $[x, x + dx]$ is $k(dx/(B - 1))((B - x)/(B - 1))^{k-1}$, corresponding to the choices for the cheapest variable and the event that variable is in $[x, x + dx]$ and all the 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 $(B + x(k - 1))/xk$. Thus, we have

$$E_k\left[\frac{Y}{X}\right] = \int_1^B \frac{k}{B-1}\left(\frac{B-x}{B-1}\right)^{k-1}\frac{1}{x}\frac{B+x(k-1)}{xk}dx$$
$$= \frac{B}{B-1}\left(\left(\int_1^B \left(\frac{B-x}{B-1}\right)^{k-1}\frac{dx}{x}\right) + \frac{k-1}{k}\right) \quad (2)$$
$$\leq \frac{B}{B-1}\left(\min\left\{\log B, \frac{B-1}{k}\right\} + \frac{k-1}{k}\right).$$

□

5.3. Unbounded Distributions. 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 the question is

$$E_k\left[\frac{Y}{X}\right] = \int_1^\infty \frac{kf(x)(1-F(x))^{k-1}}{x}dx$$
$$\times \int_x^\infty y(k-1)\frac{f(y)(1-F(y))^{k-2}}{(1-F(x))^{k-1}}dy$$
$$= k(k-1)\int_1^\infty \frac{f(x)}{x}dx\int_x^\infty yf(y)(1-F(y))^{k-2}dy. \quad (3)$$

Substituting accordingly, we immediately obtain the following corollaries for some natural distribution functions.

Corollary 5.7. *Let $NC_\lambda = \{n, r, k, F_\lambda, B\}$ with radius $r \geq \hat{r}$ and $k \leq nr^2/\log(nr^2)$ communities and F_λ the exponential distribution translated by $+1$ with parameter λ. The expected frugality ratio of VCG for NC_λ is at most $4\sqrt{2}$ w.h.p.*

Now, consider the distribution F_{recip} obtained by taking reciprocals of random variables chosen according to the uniform distribution on the unit interval $(0, 1]$.

Corollary 5.8. *Let $NC_{recip} = \{n, r, k, F_{recip}, B\}$ with radius $r \geq \hat{r}$ and $k \leq nr^2/\log(nr^2)$ communities. The expected frugality ratio of VCG for NC_{recip} is $2\sqrt{2}((k-1)/(k-2))$ w.h.p.*

In fact, for this distribution, we may say something much stronger.

Lemma 5.9. *Let $NC_{recip} = \{n, r, k, F_{recip}, B\}$ with radius $r \geq \hat{r}$ and $k \geq nr^2$ communities. The frugality ratio of VCG for NC_{recip} is at most $2e^3\sqrt{2}$ w.h.p.*

Proof. Due to the geometric bin properties, all that is needed is 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 which is $O(1/nm)$, where $m = 8/r^2$ 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, it is clear that $q = \Theta(nr^2/8)$ by coupon collection. The event that $1/X \geq e^3(1/Y)$ implies that $q - 1$ reciprocals chosen uniformly at random all lay in $(0, 1/e^3)$, the probability of which is $q/e^{3(q-1)}$. Thus, $\Pr[Y/X \geq e^3] = \Pr[1/X \geq e^3(1/Y)] < q/e^{3(q-1)}$, where X is the cheapest and Y is the second cheapest. Moreover, $q/e^{3(q-1)} \leq q/n^2 = r^2/n$ by choice of q, completing the proof. □

By noting, 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 to the above gives the following.

Lemma 5.10. *Let $NC_\lambda = \{n, r, k, F_\lambda, B\}$ with radius $r \geq \hat{r}$ and $k \geq nr^2$ communities. The frugality ratio of VCG for NC_λ is $O(1)$ w.h.p.*

6. Experimental Results

We now complement our theoretical investigations using network simulations. For this, in addition to the IAM and communities model on the random geometric graph (for which we provided theoretical bounds), we simulate several models of realistic wireless networks. While the theoretical bounds pertain to worst-case guarantees on the frugality ratio, in the simulations we can observe the distribution of frugality ratios, in many cases seeing much better performance than the worst-case bounds.

In our simulations, we consider networks consisting of 500 nodes. The primary reason for considering such a network is that while we have theoretical justification of bounds for the asymptotic case, we may obtain experimental justification for smaller numbers of nodes. This is particularly important as most realistic ad-hoc networks do not have more than 500 nodes [36]. Thus, 500 nodes is a reasonable size to consider; not so small that the FR obtained would be small trivially, something that we confirmed by simulations that showed that the FR is lower when the number of nodes in the graph is smaller. On the other hand, 500 nodes is not so large as to be redundant to the theoretical bounds and unrealistic in most circumstances.

We take the radius to be on the order to guarantee asymptotic connectivity. In particular, we take $r = 1.5\sqrt{\log n/n} = 0.16723$, where n is the number of nodes. For some distribution of the 500 nodes on the unit square, every pair of nodes within distance r of each other is connected by an edge.

As mentioned in Table 1, we ran simulations on 4 different community models.

(1) The first model has a small number of communities with a large number of nodes in each community. In this model, we choose the number of communities $k = 5$ (note that $\log n/\log\log n \approx 3$);

(2) The second model has a large number of communities with a small number of nodes in each community. Here, we choose k to be $\lfloor n/\log n \rfloor = 55$;

(3) The mixed model has a small number of large communities (4) mixed with a large number of small communities (20). With probability of 0.7, a node decides to join a large community and then chooses one of the large communities uniformly at random. With probability 0.3, a node decides to join a small community and then chooses one of the small communities uniformly at random. We choose these numbers so that the large communities will be about the same size as in the $\log n/\log\log n$ model, and the small communities will be about the same size as in the $n/\log n$ model;

(4) The individual agent model (IAM) where all nodes are independent agents.

Each community chooses a price uniformly at random $P \in [c_{\min}, c_{\min} + B]$. We distinguish the case where $(c_{\min} + B)/c_{\min}$ is small (the UAR-BDD model) from the case where $(c_{\min} + B)/c_{\min}$ is large (approximating the UAR-UNBDD model). In the UAR-BDD model, we chose $c_{\min}$ to be 1 and $B = 2$. This model is close to reality; in fact, we checked the prices in the wireless market (for a single area) and found that the ratio between one community's price to another's is less than 2.5 (our ratio can be maximum 3). In the UAR-UNBDD model, we chose $\epsilon = 10^{-5}$. The price is chosen uniformly at random from the range $[\epsilon, 1]$, this models a large variation in prices.

We ran simulations on two different graph structures, the random geometric graph $G(n, r)$ with n nodes and radius r and the randomly clustered structure G_{CL} (which we refer to simply as the clustered graph model). In G_{CL}, a realistic model when incorporating communities, each community has a territory determined by a center and territory radius r_t. We fix the same radius for all the communities, allowing each community to choose its center uniformly at random. For the model with small number of communities, we chose $r_t = 0.6$, for the rest of the models, we chose $r_t = 0.5$. We chose these numbers so that the assumption that there is no monopoly is preserved. An agent chooses its location uniformly at random within r_t of the center of its community.

We split the results into three parts: (1) UAR-BDD for $G(n, r)$, (2) UAR-UNBDD for $G(n, r)$, and (3) UAR-BDD for G_{CL}. Finally, we also discuss the sensitivity for price changes.

In general (unless we mention otherwise), the figures that will be shown below are combined from 3 different seeds, with measurements for around 75,000 (source, destination) pair samples.

6.1. The UAR-BDD Cost Model on RGG. In this section, we focus only on random geometric (RG) graphs where each community chooses uniformly at random a price $P \in [1, 3]$.

We compare the frugality ratio (FR) between the different community models. Figure 1 shows that, overall, we get a very good ratio; in fact, more than 95% of all the cases have frugality ratio below 1.4, which is much lower than the upper

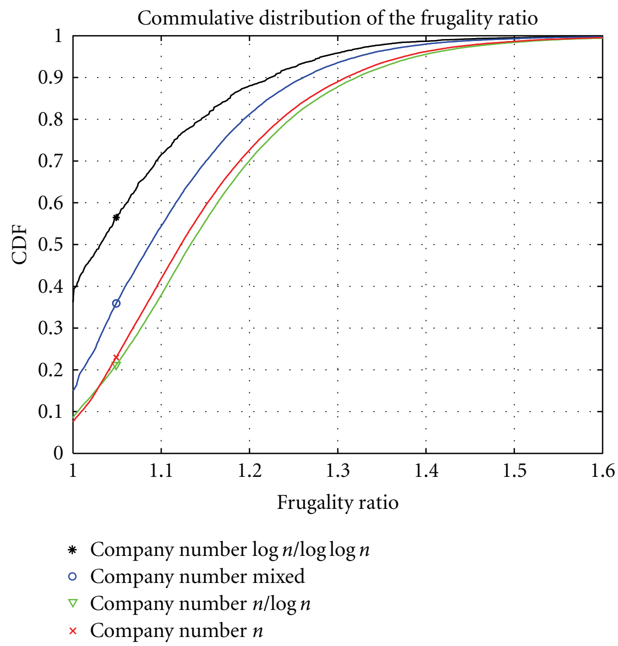

Figure 1: Comparing FR in the RGG with UAR-BDD cost distribution.

bounds that were proven in Theorem 5.3 for the IAM and Theorem 5.5 for the community models. In addition, we can see that when there are a small number of big communities in the graph, we get better results. For example, 88% of the cases in the $\log n/\log\log n$ model and 80% of the cases in the mixed models have a frugality ratio FR $\leq$ 1.2. In the $n/\log n$ model only 70% of the cases have a frugality ratio FR $\leq$ 1.2. We propose the following explanation for this. When there is a large number of small communities, any particular community will have few agents in a given area (if at all). This, combined with the range of costs being bounded, may cause the shortest path to contain many communities. When a path contains many communities, there are more overpayments to make, which leads to a higher frugality ratio. In fact, Figure 2 supports the first part of our explanation, showing that 98% of the cases for the $\log n/\log\log n$ model, and 70% for the mixed model have at most 3 communities in the shortest path compared to only 48% of the cases in the $n/\log n$ model.

We further justify the proposed explanation by showing that when there are more communities on the shortest path the frugality ratio is higher. In Figure 3, we represent each community by a box and whisker plot. The box has lines at the lower quartile, median, and upper quartile values. The whiskers are lines extending from each end of the box to show the extent of the rest of the data.

Outliers are data with values beyond the ends of the whiskers. We can see that, in general, the frugality ratio slowly increases as the number of communities on the shortest path increases (One may notice that when the number of communities is 5 and the number of communities on the shortest path is 5, the FR actually decreases; however, looking at Figure 2, we can see that we have few samples in

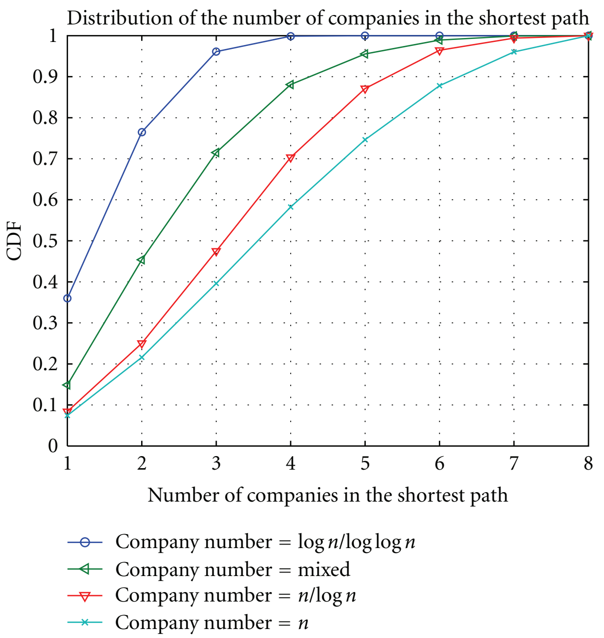

Figure 2: CDF of the number of communities in the spath (shortest path) on the RGG with UAR-BDD cost distribution.

this case since the curve almost reaches the 100%, and this is in fact an unusual case to happen). However, we can see that the *worst-case* frugality ratio is higher when the shortest path has fewer communities since the cheapest community is more dominant (a fact which we confirmed by the data).

In summary, a model with a small number of communities has lower frugality ratio. We will show next that this trend is consistent even when the costs are not bounded by a small constant, and therefore longer fluctuate in FR values may be expected.

6.2. The UAR-UNBDD Cost Model on RGG. In this subsection, we still focus on the random geometric graph, but now we will observe the UAR-UNBDD cost distribution model where a community's price $P \in [10^{-5}, 1]$. This represents a large variation in cost.

Overall, the results are consistent with the results of the previous section. Figure 4 shows that the frugality ratio in more than 99% of the cases is bounded by 4. The frugality ratio is, as expected, much larger than in the previous bounded case since the ratio between one community's price and another can be very big; however, it is still bounded by a small constatnt. Here, again, the $n/\log n$ model has the worst results overall. For example, in the $n/\log n$ model, only 83% of the cases have FR ≤ 2, whereas in the rest of the models, this happened for at least 95% of the cases. We give a similar explanation here, claiming that when all of the communities are small and there are many communities there is a higher probability to have more communities in the shortest path. Further, now the extra cost that is paid to each community can be much larger due to the price distribution. Figure 5 verifies that the likelihood of having more communities on the shortest path grows as the number of small communities grows. For example, we can see that almost 100% (resp., 70%, 40%) of the paths have up to 3 communities in the shortest path in the $\log n/\log\log n$ (resp., mixed, $n/\log n$) models.

As in the previous sections, Figure 6 shows that when the number of communities on the shortest path grows, the FR is higher. Also, we can see very clearly again that the worst-case frugality ratio is higher when the number of communities in the shortest path is smaller.

The difference here is in the individual agent model results. In the previous section, the individual agent model had similar results to the $n/\log n$ model (large number of communities); however, in the unbounded cost distribution, it behaves like the other models. The reason for this may be that agents with very low prices have less effect on paths in comparison to the $n/\log n$, in which it is likelier to have a very cheap community in every bin.

We can see that even with weaker assumptions on the cost distribution we still obtain very low frugality ratios. The results remain consistent with the previous results, and models with smaller numbers of communities still yield better results in terms of the frugality ratio.

6.3. The UAR-BDD Cost Model on a Cluster Graph. We now consider the clustered graph model with UAR-BDD cost distribution. Figure 7 shows one example of a clustered graph in the mixed model. For a better clarity, we present only 8 communities (4 large communities and 4 small communities) out of 24 communities in total.

As Figure 8 shows compared to Figure 1, there are some differences in these results. The frugality ratio of the $\log n/\log\log n$ model is not as good as it was in the random-geometric graph, for example, 70% of the clustered cases have FR ≤ 1.2 where in the random geometric graph 88% of the cases have FR ≤ 1.2. Our explanation is that in the cluster graph with small number of communities the probability of having more communities on the shortest path is higher since now the nodes of every community are located only in the community's territory and not spread all over. Figure 9 confirms our explanation. We can see in the figure that in 21% of the $\log n/\log\log n$ cases there is only one community in the shortest path, compared to the random geometric graph (Figure 2) where 37% of the cases have one community in the shortest path.

We can see another big difference in the results of the $n/\log n$ model. In the random geometric graph, the $n/\log n$ model had the worst results (Figure 1); however, in the clustered graph (Figure 8), it has the best results, for example, 81% of the cases in the clustered graph have FR ≤ 1.2 compared to 70% in the random geometric graph. In contrast to the previous case, here we have large number of small communities, so in the random geometric graph, the probability of having more communities on the shortest path increases since each community has small number of nodes that spread all over the space. On the other hand, in the clustered model, the community's nodes are located inside the community's territory, and the shortest path crosses territories where within a cheap territory it probably

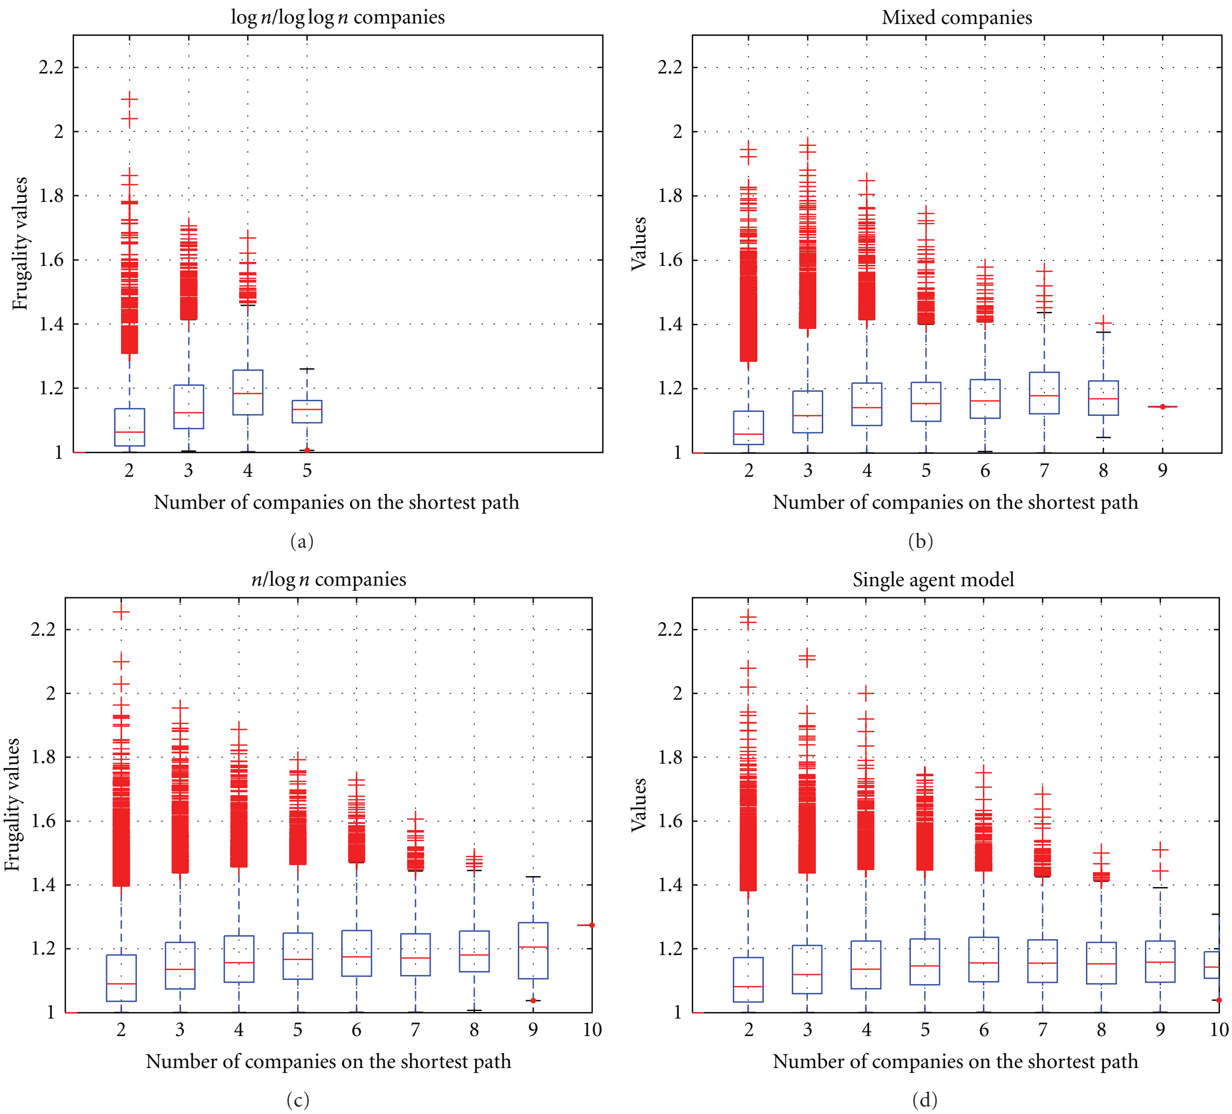

Figure 3: Frugality ratio versus number of communities in the spath, on the RGG with UAR-BDD cost distribution.

uses only its nodes. Figure 9 confirms our explanation, for example, 58% of the cases in the clustered graph has 3 communities or less in the shortest path compared to 48% in the random geometric graph.

Consistent with Sections 6.1 and 6.2, Figure 10 shows that as the number of communities in the shortest path increases, the FR increases, and the *worst-case* FR decreases.

In summary of this subsection, we can see that there are some major differences between the random geometric graph results and the clustered graph results. In particular, a model with large number of small community in clustered graph has better results in terms of frugality ratio.

In addition, we checked the sensitivity of payments when the prices changed by individual community, and the other communities' costs remain the same. We ran two different sets of simulations one on the random geometric graphs and another on the clustered graphs both using the UAR-UNBDD cost model. The simulations show that as expected from a truthful mechanism, a community gets about the same revenue if it has the cheapest price, and it does not matter what the price is. Overall, a community has a lower probability to be affected by price changes in the clustered model. However, once it is effected, the effect will be stronger in the cluster model. We omit the simulation results due to space constraints.

7. Hardness of Extensions

As has been noted, both simulation results and related work on the traditional path auction model [8–10, 30] suggest that a mechanism that minimizes some weighting of total path costs by the number of communities on the path may have a lower frugality ratio than VCG. Another immediate question is how one might generalize the $\sqrt{n}$ mechanism [10] which is known to be up to $\sqrt{n}$ times more frugal than VCG for the traditional model to our community model. In the context

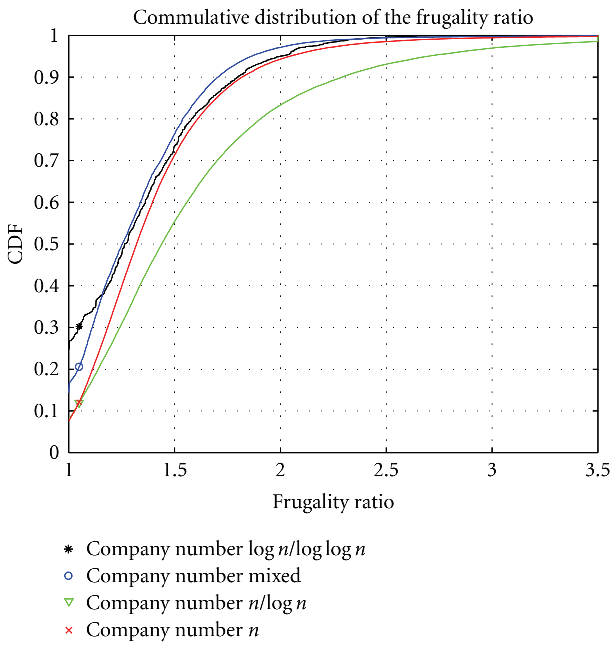

FIGURE 4: Comparing frugality ratio in the RGG with UAR-UNBDD cost distribution.

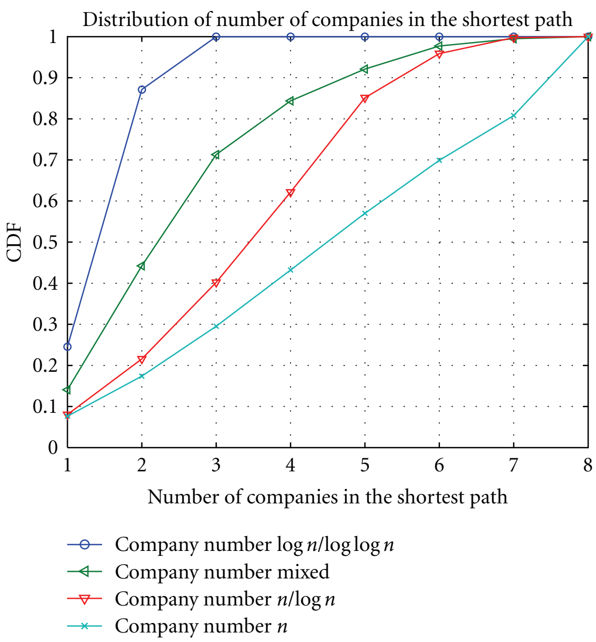

FIGURE 5: CDF of the number of communities in the spath on the RGG with UAR-UNBDD cost distribution.

of extensions and variations of truthful mechanisms for the community model, it is further worthwhile to investigate VCG for the community model under other cost models. Unfortunately, we have found that many approaches in these general directions turn out to be NP-complete, some even strongly approximation hard.

7.1. NP-Hardness of Extending $\sqrt{n}$. The first step of the $\sqrt{n}$ mechanism is to find the least cost edge-disjoint cycle through s and t, and then the winner of each vertex disjoint partition of the cycle is decided according to a function weighting the cost of each side of the partition by square-root of the number of edges on that side. In terms of the community model, this would correspond to finding at least some community disjoint cycle through s and t (It should be noted that two community disjoint paths are not necessary for the no-monopoly condition. E.g., consider $k = 3$ and a graph consisting of three length-two 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, as we shall see that results apply to both cases) according to the communities they belong to. Finding a community disjoint cycle is the same as finding a color-disjoint cycle. We can show that this problem is NP-complete by a reduction from 3-SAT. A similar problem has independently been shown to be NP-complete in [37] as well.

Lemma 7.1. *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.*

7.2. APX-Hardness of Natural Extensions. 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. 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 positive number of community i's nodes, which may be a more natural model for some service providers). Our reduction is an approximation preserving reduction from the minimum monotone satifying assignment (MMSA_3) problem, which is known to be $2^{\log^{1-o(1)} n}$ hard to approximate [38, 39]. While there are closely related approximation hardness results under various names [40, 41], our result and reduction are both more general and more direct.

First, for ease of notation, let us note the following:

$$\forall 0 < x < 1, \quad 2^{\log^{1-o(1)} n} > n^x. \tag{4}$$

Now, define a natural class of truthful mechanisms for path auctions in the community model.

Definition 7.2. We call a truthful mechanism for path auctions in the community model (with per unit costs) a *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 nondecreasing function g, where q is the number of communities on P and p is the total cost of P.

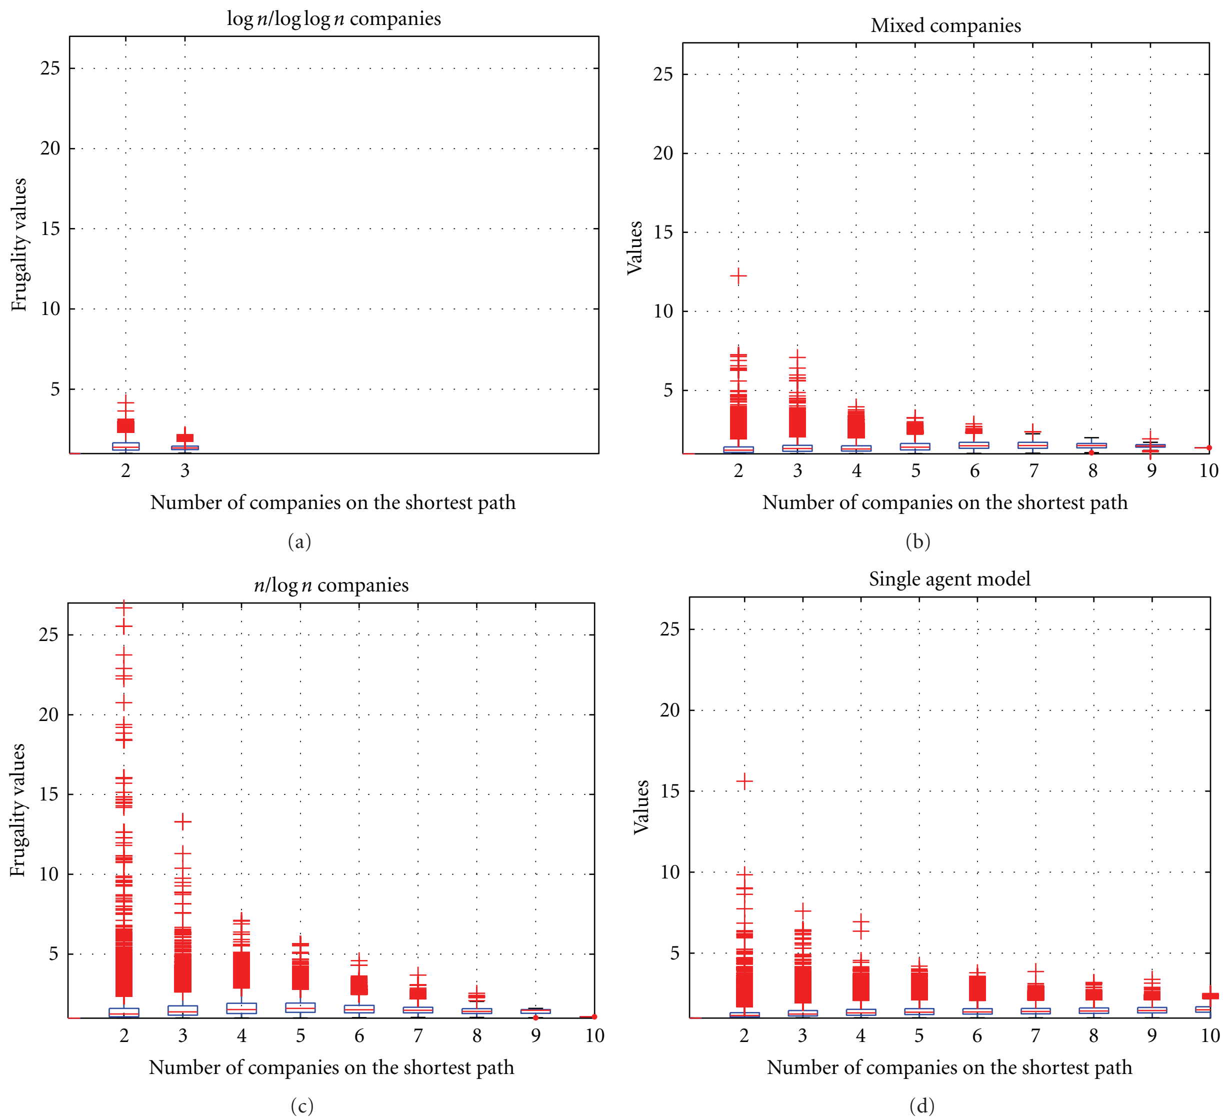

Figure 6: Frugality versus number of communities in the spath, with UAR-UNBDD cost distribution, in the RGG.

In particular, denote such a mechanism as a (f, g) min-agent mechanism.

Now, we proceed to our hardness results.

Theorem 7.3. *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.*

Proof. Let σ be an instance of MMSA_3, and let $\mathfrak{C}$ be the corresponding circuit. Let $C_1, C_2, \ldots, C_z$ be the highest level (3rd level) ANDed clauses. For each i, let $E_{i,1}, E_{i,2}, \ldots, E_{i,q_i}$ be the (2nd level) corresponding ORed clauses within. Finally, for each i, j, let $x_{i,j,1}, x_{i,j,2}, \ldots, x_{i,j,p_{i,j}}$ be the (1st level) corresponding ANDed variables within. The entire structure has a straightforward representation as a SERIES-PARALLEL-SERIES (SPS) graph with variables corresponding to colored nodes in the most internal SERIES layers. Thus, we convert into an instance of the min-agent problem as follows. Let $\widehat{G}$ be the SPS graph of σ with C_1 connected to s and C_z connected to t, and each color corresponding to a community. Note that adding copies of the same variable within the innermost layer does not change the solution of the MMSA_3 problem in any way, not even approximately. Set cost function c to be simply a unit cost function. Now, what we want is for all the $E_{i,j}$s for fixed i to become the same length within C_i so that $g(p)$ is the same for all paths from s to t. So, let $L_i = \max_j |E_{i,j}|$, and for each $E_{i,j}$ choose some variable $x_{i,j} = x_{i,j,b}$ and make $L_i - |E_{i,j}|$ copies of $x_{i,j}$ in $E_{i,j}$. We have that $p = \sum_{i=1}^{z} L_i$ for all $s - t$ paths. Feed the resulting graph G into a (f, g) min-agent mechanism. It is clear that if the path selected is α factor of the optimum path with respect to minimization function $f(k)g(p)$, then since the optimum path also has

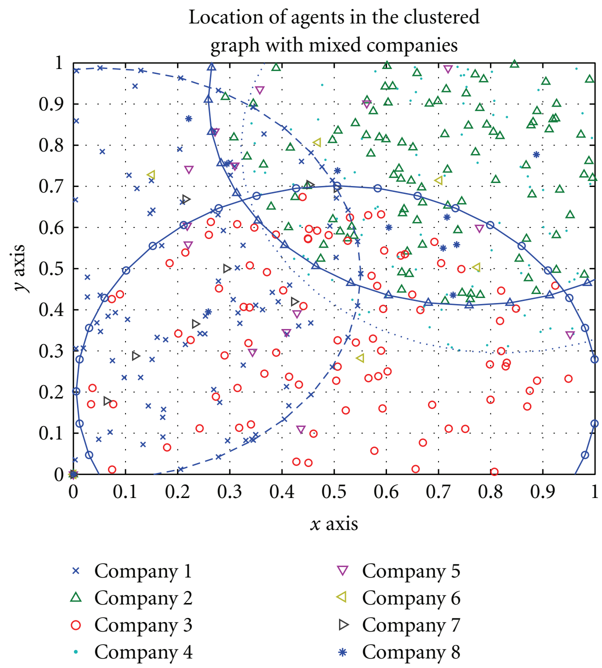

Figure 7: Example for a clustered graph in the mixed model, the figure shows only 8 communities (out of 24 communities in total).

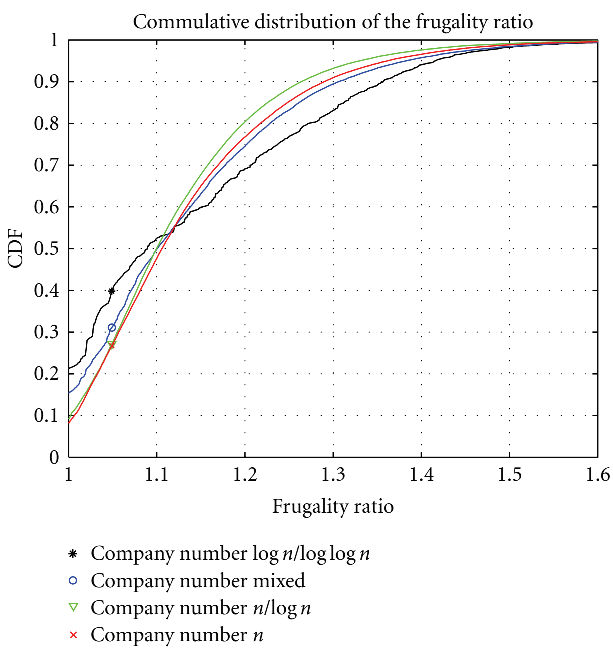

Figure 8: Comparing frugality ratio in the cluster graph with UAR-BDD cost distribution.

the same $g(p)$, the path selected must be within α factor of optimal $f(k)$, thus solving $MMSA_3$ to within an $f^{-1}(\alpha)$ factor approximation as well. Thus, the theorem follows from Remark 2 and approximation hardness of $MMSA_3$. $\square$

Taking g to be a constant function, we may obtain the following corollary.

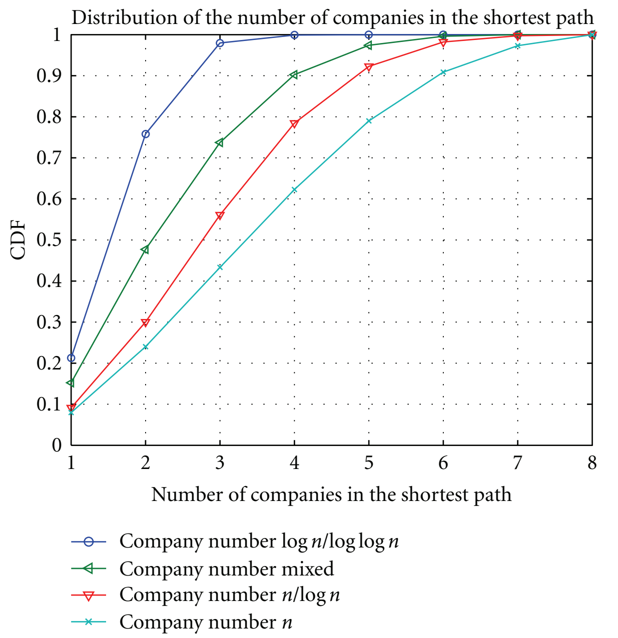

Figure 9: CDF of the number of communities in the spath with UAR-BDD cost distribution on a cluster graph.

Corollary 7.4. *VCG for the community model under fixed community subnetwork entrance fees is hard to approximate to within* k^x, *for any* $0 < x < 1$, *given k total communities.*

8. Discussion

8.1. Summary. We have considered the cost of the generalized second price path auction (VCG) for incentive-compatible source-destination routing in a random ad-hoc network setting in which nodes may potentially be grouped, where we refer to each grouping as a *community*. We have proven bounds on the *frugality ratio* in this setting, namely, the ratio of the total payment made to every independent agent in order to ensure truthfulness over the actual cost of the shortest path. Whereas it is well known that this ratio may be arbitrarily bad for VCG as well as for truthful mechanisms in general, even when every edge or node is of unit cost, we are motivated by the understanding that worst-case results are often exhibited on pathological rather than typical case graphs and cost distributions. Therefore, we have asked our questions on a model of the *typical* case class of graphs for ad-hoc networks, namely, random geometric graphs. And, we have considered both arbitrary costs without any assumption made on the distribution, as well as costs drawn from natural probability distributions. In all such cases, we have shown that VCG extended to capture communities exhibits constant frugality ratio with high probability given some reasonable assumptions on the maximum cost offset as well as certain natural unbounded cost distributions, and when no assumption is made on the maximum cost offset, the performance of VCG with communities is still very efficient (logarithmic in the offset) in expectation. Simulation experiments demonstrate even

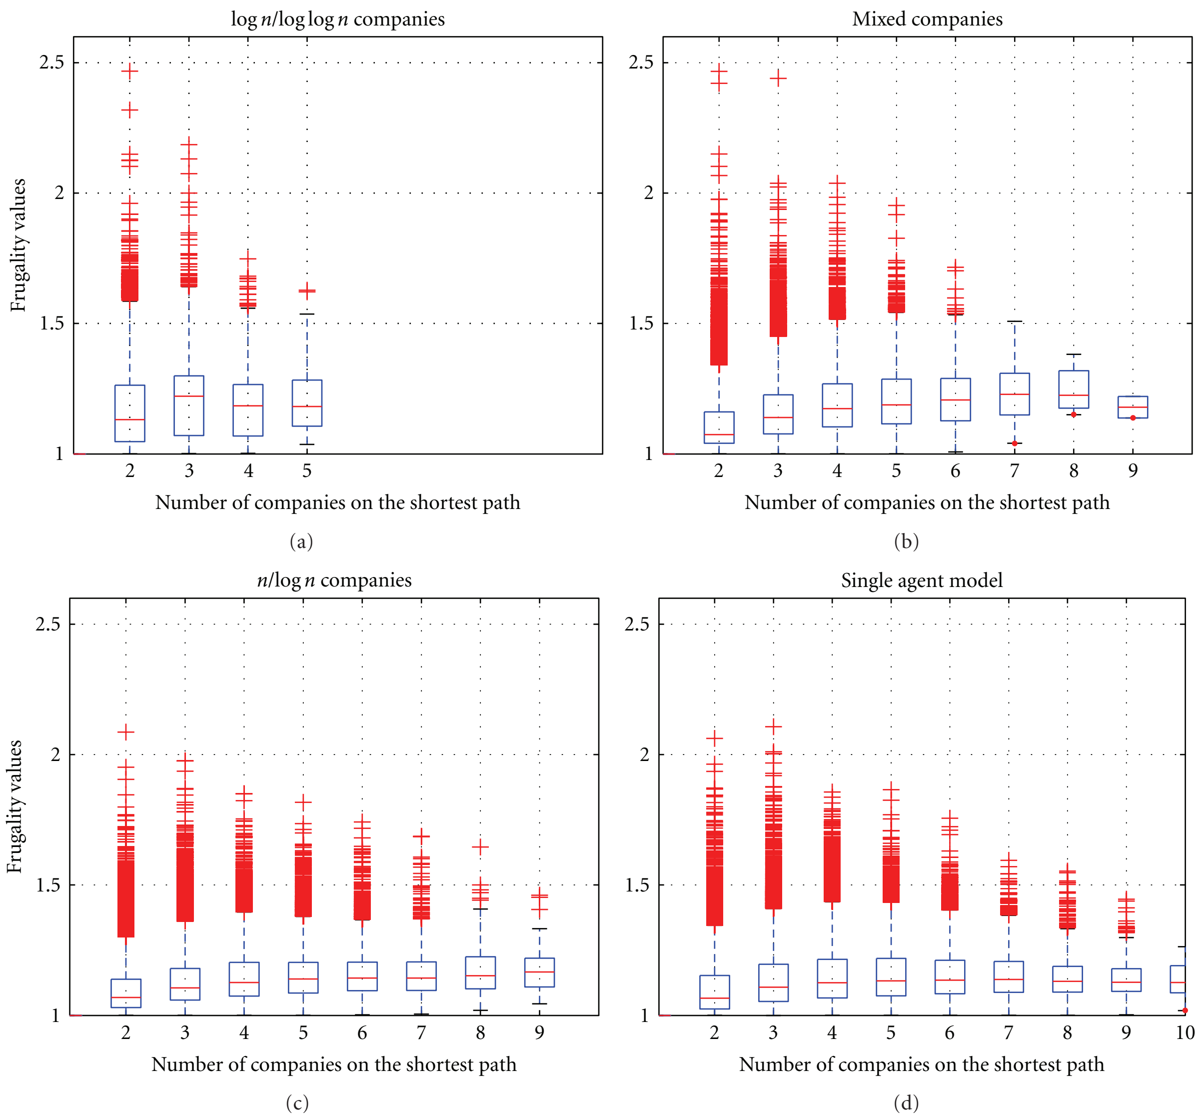

Figure 10: Frugality versus number of communities in the spath, with UAR-BDD cost distribution, in the clustered graph.

smaller frugality bounds than those given via theoretical results, and simulations also demonstrate small constant frugality ratio in the clustered community setting which was not considered in the theoretical bounds. We are the first to prove bounds on frugality for random geometric graphs. We are also the first to consider path auctions in the community setting, as well as the first to consider frugality bounds under natural cost distributions. None of our bounds follow from previous work and, thus, indeed pose as pleasant surprises.

Despite the positive note obtained from low frugality ratio in many natural settings for random ad-hoc networks, experimental results do suggest that frugality would still be improved had the truthful mechanism taken the *number* of communities on the chosen path (and thus the number of overpayments made) into account in addition to the total cost of the path which it already considers. Therefore, we examined general ways in which a truthful mechanism may take into account such considerations, but in all cases came across hurdles in computational complexity. We proved that for such classes of truthful mechanisms, even approximately computing the winning path is provably hard theoretically. Therefore, we contend with the positive note on low frugality ratio for VCG extended for communities, which we see is not only good but also “as good as it gets.”

8.2. Robustness of Results. We would like to note the robustness of our results under some alternative modeling assumptions. First, the assumption that all nodes in a community have the same cost is in a way strongly defining the community concept. If this assumption was wanting of relaxation, then note that one may take the bounds we have proven for the individual agent models (in which each node is its own community) as an extreme worst case, or simply enlarge the number of communities (change parameter k)

considered to obtain tighter bounds. Second, if we were to relax the assumption of bounded broadcast radii, but rather assume a completely connected graph with only the cost distribution affecting payments, then this corresponds to the special case of taking the radius $r = \sqrt{2}$, and correspondingly low frugality bounds follow from such one (as in fact taking larger radius only relaxes the hypothesis statement for the bounds given). Lastly, we would like to note the benefit of generality in our theoretical results and experimental modeling. The results do not depend on any particular application or network type save that the fundamental communication medium is wireless. Similarly, the community concept is independent of application and may thus apply to any setting from providers in a cellular network to nodes grouped together based on trust or similarity in a sensor network. As costs too are kept general, nodes may dictate their own valuations based on any number of objectives they wish to optimize (e.g., broadcast energy, remaining power, the general desire or lack thereof to participate in the transfer of information from particular source-destination pairs, etc.). The only assumption required for the employment of game-theoretic principles, as we have considered in this work, is the existence of autonomous agents in the network. Like much recent work in this selfish autonomous network setting, while the assumption of true autonomy may yet be suspect in most current applications, it is a state of affairs that we must be prepared for.

Appendix

Proof of Theorem 5.3. Let s and t be an arbitrary source and sink pair, and let $\text{SP} = \langle v_0, v_1, v_2, \dots, v_d \rangle$ with $s = v_0$, $t = v_d$, $|\text{SP}| = d$ be a shortest path between s and t. By Lemma 5.1 and the choice of r, every $(r/2) \times (r/2)$ square region has $\Theta(nr^2)$ nodes. So, for each i, let l_i be the line between v_{i-1} and v_{i+1}. Partition l_i into $d(v_{i-1}, v_{i+1})/(r/2)$ intervals of length $r/2$ each. Depending on how close l_i is to a boundary of the unit square, it is clear that we may partition part of the area between v_{i-1} and v_{i+1} with at most four diagonally adjacent $(r/2\sqrt{2}) \times (r/2\sqrt{2})$ squares S_1, S_2, S_3, S_4 as in Figure 11. By Lemma 5.1 and the choice of r, each S_j has $\Theta(nr^2/8)$ nodes w.h.p. Moreover, each $S_j \cup S_{j+1}$ forms a clique. Now, partition the range $R = [c_{\min}, c_{\min} + B]$ into $b = (nr^2/8)/2\log(nr^2/8) \leq nr^2/4\log(nr^2)$ intervals of length B/b each. Since each node picks a cost independently and uniformly at random from range R, by Lemma 5.1 and the choice of b, w.h.p. for each S_j, for each interval $q \in \{1, \dots, b\}$, there exist at least $\log|S_j| \geq \log\log(nr^2/8)$ nodes in S_j with cost in the first interval $[c_{\min}, c_{\min} + B/b]$. Since the probability that any two nodes pick the exact same cost from a continuous range of costs is zero, then we have that w.h.p., the second cheapest community for each S_j has cost at most $c_{\min} + B/b$, w.h.p. Choosing such nodes u_2 and u_3 in S_2 and S_3, respectively, and connecting edges (v_{i-1}, u_2), (u_2, u_3), and (u_3, v_{i+1}) by the clique property, we have that the cost of the cheapest path between s and t without node v_i is at most $2((B + c_{\min})/bc_{\min})$ more than the cost of SP with node v_i. Thus, the lemma follows. □

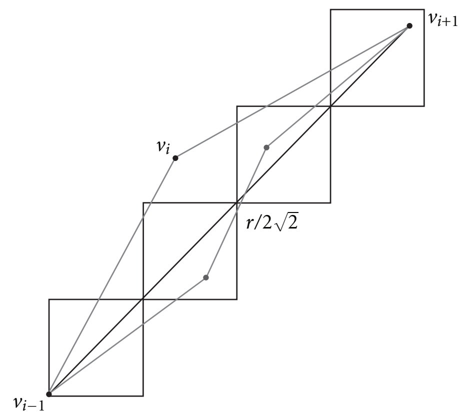

Figure 11: IAM and alternate paths.

Proof of Theorems 5.2 and 5.4. Let s and t be an arbitrary source and sink pair, and let $\text{SP} = \langle v_0, v_1, v_2, \dots, v_d \rangle$ with $s = v_0, t = v_k, |\text{SP}| = d$ be a shortest path between s and t. By Lemma 5.1 and the choice of r, every $(r/2) \times (r/2)$ square region has $\Theta(nr^2)$ nodes. So, for any pair of points v_i and v_{i+1}, since $d(v_i, v_{i+1}) \leq r$ w.h.p., there exists another point $\overline{v_i}$ such that $\overline{v_i} \neq v_j$ for j and $d(v_i, \overline{v_i}) \leq r$ and $d(\overline{v_i}, v_{i+1}) \leq r$. Therefore, the VCG payment in the edge-agent model to edge $e_i = (v_i, v_{i+1})$ is at most $2(c_{\min} + B)$, whereas the actual cost of e_i is at least $c_{\min}$.

Regarding the node-agent IAM, for each v_i, partition the line between v_{i-1} and v_{i+1} as in the proof of Theorem 5.3. Proceeding similarly, it is easy to see that due to the choice of r, the worst-case frugality ratio in the node agent model is $2(1 + B/c_{\min})$.

Regarding the community model in either case, partition according to the proof of Theorem 5.5. Then, for each A_{j_i}, we obtain a path of cost at most $2\sqrt{2}(d(v_{j_i,0}, v_{j_i,f})/r)c_{\min} + B$ 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 frugality ratio is at most $2\sqrt{2}(1 + B/c_{\min})$. Summing over each L_{j_i}, we just obtain the same ratio. This characterizes the payment to community j. Moreover, clearly, the argument is the same for any community since the scaling by distance is lost. Thus, the lemma follows. □

Acknowledgments

The authors 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, as well as the anonymous reviewers for very useful comments. This research was sponsored in part by the NSF Grants CCF-0427202 and CCF-0546170.

References

[1] C. H. Papadimitriou, "Algorithms, games, and the internet," in *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC '01)*, pp. 749–752, ACM, Crete, Greece, July 2001.

[2] L. Anderegg and S. Eidenbenz, “Ad hoc-VCG: a truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents,” in *Proceedings of the 9th Annual International Conference on Mobile Computing and Networking (MobiCom ’03)*, pp. 245–259, ACM, San Diego, Calif, USA, September 2003.

[3] D. Kreps, *A Course in Microeconomic Theory*, Princeton University Press, Princeton, NJ, USA, 1990.

[4] N. Nisan and A. Ronen, “Algorithmic mechanism design,” *Games and Economic Behavior*, vol. 35, no. 1-2, pp. 166–196, 2001.

[5] W. Vickrey, “Counterspeculation, auctions, and competitive sealed tenders,” *Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.

[6] E. H. Clarke, “Multipart pricing of public goods,” *Public Choice*, vol. 11, no. 1, pp. 17–33, 1971.

[7] T. Groves, “Incentives in teams,” *Econometrica*, vol. 41, pp. 617–631, 1973.

[8] A. Archer and E. Tardos, “Frugal path mechanisms,” in *Proceedings of the 13th Annual ACM/SIAM Symposium on Discrete Algorithms (SODA ’02)*, pp. 991–999, San Francisco, Calif, USA, January 2002.

[9] K. Talwar, “The price of truth: frugality in truthful mechanisms,” in *Proceedings of the the 20th Annual Symposium on Theoretical Aspects of Computer Science (STACS ’03)*, vol. 2607 of *Lecture Notes in Computer Science*, pp. 608–619, Berlin, Germany, 2003.

[10] A. R. Karlin, D. Kempe, and T. Tamir, “Beyond VCG: frugality of truthful mechanisms,” in *Proceedings of the 46th IEEE Annual Symposium on Foundations of Computer Science (FOCS ’05)*, pp. 615–624, Las Vegas, Nev, USA, October 2005.

[11] H. Sun and J. Song, “Strategy proof trust management in wireless ad hoc network,” in *Proceedings of the Canadian Conference on Electrical and Computer Engineering*, pp. 1593–1596, Ontario, Canada, May 2004.

[12] W. Wang, X. Y. Li, S. Eidenbenz, and Y. Wang, “OURS: optimal unicast routing systems in non-cooperative wireless networks,” in *Proceedings of the 12th Annual International Conference on Mobile Computing and Networking (MobiCom ’06)*, pp. 402–413, ACM Press, Los Angeles, Calif, USA, September 2006.

[13] S. Zhong, L.E. Li, Y. G. Liu, and Y. R. Yang, “On designing incentive-compatible routing and forwarding protocols in wireless ad-hoc networks: an integrated approach using game theoretical and cryptographic techniques,” in *Proceedings of the 11th Annual International Conference on Mobile Computing and Networking (MobiCom ’05)*, pp. 117–131, Montréal, Canada, September 2005.

[14] M. Musolesi and C. Mascolo, “A community based mobility model for ad hoc network research,” in *Proceedings of the 2nd International Workshop on Multi-Hop Ad Hoc Networks: From Theory to Reality (REALMAN ’06)*, pp. 31–38, ACM Press, Florence, Italy, May 2006.

[15] G. Kortuem and Z. Segall, “ Wearable communities: augmenting social networks with wearable computers,” *IEEE Pervasive Computing*, vol. 2, no. 1, pp. 71–78, 2003.

[16] S. Schulz, K. Herrmann, R. Kalckloesch, and T. Schwotzer, “Towards trust-based knowledge management in mobile communities,” in *Proceedings of the AAAI Symposium on Agent-Mediated Knowledge Management*, 2003.

[17] M. Musolesi, S. Hailes, and C. Mascolo, “An ad hoc mobility model founded on social network theory,” in *Proceedings of the 7th ACM International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM ’04)*, pp. 20–24, ACM Press, Venice, Italy, October 2004.

[18] A. Patrick, C. Olivier, P. Jean-Marc, J. Bernard, M. Ludovic, and P. Ricardo, “Security in ad hoc networks: a general intrusion detection architecture enhancing trust based approaches,” in *Proceedings of the 1st International Workshop on Wireless Information Systems in Conjunction with the 4th International Conference on Enterprise Information Systems*, Ciudad Real, Spain, April 2002.

[19] M. D. Penrose, *Random Geometric Graphs*, Oxford University Press, 2003.

[20] P. Gupta and P. R. Kumar, “The capacity of wireless networks,” *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388–404, 2000.

[21] A. Goel, S. Rai, and B. Krishnamachari, “Monotone properties of random geometric graphs have sharp thresholds,” *Annals of Applied Probability*, vol. 15, no. 4, pp. 2535–2552, 2005.

[22] S. Muthukrishnan and G. Pandurangan, “The bin-covering technique for thresholding random geometric graph properties,” in *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA ’05)*, pp. 989–998, Vancouver, Canada, January 2005.

[23] C. Avin and G. Ercal, “On the cover time of random geometric graphs,” in *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP ’05)*, pp. 677–689, Springer, Lisbon, Portugal, July 2005.

[24] J. Diaz, J. Petit, and M. J. Serna, “Faulty random geometric networks,” *Parallel Processing Letters*, vol. 10, no. 4, pp. 343–357, 2000.

[25] U. C. Kozat and L. Tassiulas, “Throughput scalability of wireless hybrid networks over a random geometric graph,” *Wireless Networks*, vol. 11, no. 4, pp. 435–449, 2005.

[26] R. Motwani and P. Raghavan, *Randomized Algorithms*, Cambridge University Press, Cambridge, UK, 1995.

[27] G. Ercal, R. Izhak-Ratzin, R. Majumdar, and A. Meyerson, “Frugal routing on wireless ad-hoc networks,” in *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT ’08)*, pp. 133–144, Paderborn, Germany, 2008.

[28] L. Buttyán and J. Hubaux, “Nuglets: a virtual currency to stimulate cooperation in self-organized ad hoc networks,” Tech. Rep. DSC/2001/00, 2001.

[29] N. Nisan and A. Ronen, “Computationally feasible VCG mechanisms,” in *Proceedings of the 2nd ACM Conference on Electronic Commerce (EC ’00)*, pp. 242–252, Minneapolis, Minn, USA, October 2000.

[30] E. Elkind, A. Sahai, and K. Steiglitz, “Frugality in path auctions,” in *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA ’04)*, pp. 701–709, New Orleans, La, USA, January 2004.

[31] Y. Du, R. Sami, and Y. Shi, “Path auction games when an agent can own multiple edges,” in *Proceedings of the Workshop on the Economics of Networks, Systems and Computation (NetEcon ’06)*, Ann Arbor, Mich, USA, June 2006.

[32] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker, “A BGP-based mechanism for lowest-cost routing,” in *Proceedings of the 21st Annual ACM Symposium on Principles of Distributed Computing (PODC ’02)*, pp. 173–182, Monterey, Calif, USA, July 2002.

[33] M. Mihail, C. Papadimitriou, and A. Saberi, “On certain connectivity properties of the internet topology,” in *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS ’03)*, pp. 28–35, Cambridge, Mich, USA, October 2003.

[34] D. Karger and E. Nikolova, "VCG overpayment in random graphs," in *Proceedings of the DIMACS Workshop on Computational Issues in Auction Design*, October 2004.

[35] N. Immorlica, D. Karger, E. Nikolova, and R. Sami, "First-price path auctions," in *Proceedings of the 6th ACM Conference on Electronic Commerce (EC '05)*, pp. 203–212, Vancouver, Canada, June 2005.

[36] P. Gupta, R. Gray, and P. Kumar, "An experimental scaling law for ad hoc networks," Tech. Rep., 2001.

[37] S. Yuan and J. P. Jue, "Dynamic lightpath protection in WDM mesh networks under wavelength-continuity and risk-disjoint constraints," *Computer Networks and ISDN Systems*, vol. 48, no. 2, pp. 91–112, 2005.

[38] I. Dinur and S. Safra, "On the hardness of approximating label-cover," *Information Processing Letters*, vol. 89, no. 5, pp. 247–254, 2004.

[39] M. Alekhnovich, S. R. Buss, S. Moran, and T. Pitassi, "Minimum propositional proof length is NP-hard to linearly approximate," in *Proceedings of the 23rd International Symposium on Mathematical Foundations of Computer Science (MFCS '98)*, pp. 176–184, Brno, Czech Republic, August 1998.

[40] R. D. Carr, S. Doddi, G. Konjevod, and M. Marathe, "On the red-blue set cover problem," in *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '00)*, pp. 345–353, San Francisco, Calif, USA, January 2000.

[41] H. C. Wirth, *Multicriteria approximation of network design and network up-grade problems*, Ph.D. thesis, 2001.