

VCG Overpayment in Random Graphs

David R. Karger, Evdokia Nikolova
MIT Computer Science and Artificial Intelligence Lab
{karger, enikolova}@csail.mit.edu

October 5, 2004

Abstract

Motivated by the increasing need to price networks, we study the prices resulting from of a variant of the VCG mechanism, specifically defined for networks [11]. This VCG mechanism is the unique efficient and strategyproof mechanism, however it is not budget-balanced and in fact it is known to result in arbitrarily bad overpayments for some graphs [1, 8]. In contrast, we study more common types of graphs and show that the VCG overpayment is not too high, so it is still an attractive pricing candidate. We prove that the average overpayment in Erdős-Renyi random graphs with unit costs is $p/(2-p)$ for any n , when the average degree is higher than a given threshold. Our simulations show that the overpayment is greater than $p/(2-p)$ below this threshold, hence, together with the constant upper bound from Mihail *et al.* [12], the overpayment is constant regardless of graph size. We then find empirically that the complete graph with random edge costs has worse overpayments, higher than $\Theta((\log n)^{1.5})$. Again from simulations, we see that the power-law graph with unit costs has overpayments that decrease with graph size and finally, the power-law graph with uniformly random costs has a small constant overpayment.

1 Introduction

With the rapid growth of the Internet, there is an increasing need to automatize the labor-intensive policy setting for inter-domain routing. This can be resolved by appropriate pricing mechanisms. Integration of prices with routing protocols may also offer successful models for congestion control. Feigenbaum *et al.* [11] offer a mechanism for lowest-cost routing which integrates well with the current routing protocol. Their mechanism falls in the class of the celebrated Vickrey-Clarke-Groves (VCG) mechanisms. As such, it may require payments much larger than the true routing costs [1, 8].

Motivated by the idea of integrating the VCG mechanism into the Internet inter-domain routing graph, we turn to studying the effect of graph topology on the VCG payments. Our work is independent of the Internet routing graph per se and can be applied to any equivalent network. It can also be applied to a more general setting of interaction among multiple agents such as task allocation to teams of selfish agents.

1.1 Our Results.

An important part of our work is to formally define the concept of VCG overpayment in path auctions and to distinguish the notions of edge bonus, path bonus and average overpayment.

We study three main types of graphs: Erdős-Renyi random graphs, in which each edge is present with probability p ; complete graphs and power-law graphs. We prove that the average overpayment in Erdős-Renyi random graphs with unit costs is $p/(2-p)$ for any n , when the average degree is higher than a given threshold. Our simulations show that the overpayment is greater than $p/(2-p)$ below this threshold, hence together with the constant upper bound from Mihail *et al.* [12], the overpayment is constant regardless of graph size. A corollary of our result is that, contrary to intuition, overpayments become larger as the graph becomes denser and more competitive.

We then find empirically that again, contrary to intuition, the complete graph with uniformly random costs, which is perfectly competitive, has much worse overpayments, higher than order of $(\log n)^{1.5}$.

Lastly, our simulation results for the preferential attachment model show that power-law graphs with unit costs have small overpayment which decreases with graph size, and those with uniformly random costs have a constant overpayment.

Tangmunarunkit *et al.* [15] show that network topology generators based on degree distribution model the Internet graph much more closely than do structural generators. With this, our results suggest that the VCG mechanism, when implemented on real world networks, will result in very reasonable prices, contrary to the numerous criticisms it has received about requiring arbitrarily high payments.

1.2 Related Work.

From the Computer Science literature, Archer and Tárdoš [1] were perhaps the first to delve into how bad VCG payments can get. In fact they showed that a more general class of truth-telling mechanisms is doomed to require arbitrarily high payments and Elkind *et al.* [8] strengthened this result to include all truth-telling mechanisms. Recently, Elkind [9] also noted that payments can be reduced by removing edges from a given graph, however it is NP-hard to decide which is the optimal set of edges to remove in order to obtain the lowest possible VCG payments.

On the other hand, Feigenbaum *et al.* [11] computed the average VCG overpayment (over all source-sink destination pairs and edges on the shortest paths) on the Internet inter-domain routing graph to be 44% of the cost when all edges have unit cost. This overpayment is intriguingly low, considering the size of the graph and the results of the above papers. The authors suggested that this low overpayment may be due to the fact that the Internet is very well connected and hence quite competitive.

In an unrelated study, Chung and Lu [6] showed that the average distance in a power-law graph such as the Internet (with power-law exponent between 2 and 3) is very low: $O(\log \log n)$, where n is the number of vertices in the graph. Shortly afterwards, while studying VCG overpayments, Mihail *et al.* [12] noted that as a corollary of the average distance result, the average VCG overpayment per unit cost is also upper-bounded by $O(\log \log n)$ in the corresponding power-law model with unit edge costs. In the same work, the authors proved that the average VCG overpayment in Erdős-Renyi random graphs is between $\Omega(1/np)$ and $O(1)$. They also stated a conjecture that the overpayment is exactly $\Theta(1/np)$, *i.e.*, the overpayment decreases as the graph gets bigger and more dense. A second conjecture in the same paper was that the overpayment in power-law graphs is constant.

Note that our results disprove the first conjecture about Erdős-Renyi graphs, while our empirical findings confirm the second conjecture and give rise to a stronger conjecture, that in power-law graphs with unit costs, the overpayment decreases with the size of the graph.

2 Definitions and the VCG Mechanism

We describe the statement of the problem after Feigenbaum *et al.* [11]. We are given a network with n nodes and a set of bidirectional links between the nodes. We are interested in selecting the cheapest possible route between a source and a destination node. Since this problem is trivial when there is only one possible route, we will only concern ourselves with the biconnected components of graphs, namely the components in which there are at least two routes between any source and destination. Note that from a pricing point of view, we can only draw meaningful conclusions for biconnected networks since a monopolist can charge infinite prices.

Each edge declares a transit cost which is private information to that edge. The cost of a route is the sum of the costs of its edges. A pricing mechanism consists of possible cost reports or bids for each edge, an allocation rule (say, the cheapest possible path from a source to a destination) and a payment rule (how much we pay each edge on or off the chosen path). In general, it is hard to design a mechanism that would minimize payment to the edges and would guarantee low cost to consumers. This is why we focus on efficient mechanisms, which select the lowest-cost paths and thus minimize the system cost. In addition, we would like to have a strategyproof mechanism, that is a mechanism in which it is optimal for each edge to reveal its true cost, regardless of what other edges do.

It turns out that there is a unique strategyproof and efficient mechanism, which falls in the family of the Vickrey-Clarke-Groves mechanisms.

For a given source s and destination t , denote by $LCP(G, s, t)$ the cost of the lowest cost path from s to t in the graph G .

Theorem 1 [11] *When routing picks lowest-cost paths, and the network is biconnected, there is a unique strategyproof pricing mechanism that gives no payment to edges that carry no transit traffic. The payment to each edge e on the lowest cost path from s to t is given by*

$$v(e, s, t) = \text{cost}(e) + LCP(G - e, s, t) - LCP(G, s, t). \quad (1)$$

From here on we refer to this mechanism as the VCG mechanism. $(G - e)$ above stands for the graph G with edge e deleted. Note that each edge on a winning path is paid its cost plus an extra term $LCP(G - e, s, t) - LCP(G, s, t)$, which we will refer to as *bonus*. The bonus payment to an edge is precisely the marginal cost reduction of a route from adding the edge to the network. Another interpretation by Elkind [9] is that $v(e, s, t)$ is the threshold bid of an edge e , i.e. the highest possible bid that the edge may announce and still be on the winning lowest-cost path from s to t .

Definition 1 *The edge bonus is given by $LCP(G - e, s, t) - LCP(G, s, t)$, for every edge e on the lowest cost path from s to t . The path bonus is the sum of the bonuses of all edges on the path.*

Next, we define the notion of VCG overpayment. In all graphs we simulate uniform traffic, i.e., we have a unit flow between each pair of nodes. To be precise, assume a unit flow between each pair in one direction only, although this will not make a difference. Let V be the set of vertices in the graph G .

Definition 2 *The total cost of traffic is the sum of costs of all least cost paths,*

$$TC = \sum_{s \neq t \in V} LCP(G, s, t).$$

The total payment is the sum of all payments,

$$TP = \sum_{s \neq t \in V} \sum_{e \in LCP(G,s,t)} v(e, s, t).$$

The total bonus is the difference of total payment and total cost,

$$TB = TP - TC.$$

We may want to study several metrics—the average overpayment (and payment) per path, per unit cost, or per edge. In the case of unit costs, the last two metrics are of course the same. We define the average overpayment per unit cost, which we will just refer to as average overpayment, to be the ratio of total bonus to total cost.

Definition 3 The average VCG overpayment in a graph G with uniform traffic between all pairs of nodes, is given by

$$\text{avg overpayment} = \frac{TP - TC}{TC}.$$

The average VCG payment is the ratio of total payment to total cost, $\frac{TP}{TC} = 1 + \text{avg overpayment}$.

Note that there are alternative ways to measure the average overpayment. For example, we may divide for each edge the bonus over its cost, and then average over all such edges. However, with this approach it is not clear if we should first sum bonuses for a single edge which participates in several LCPs or if we should sum bonuses over each LCP, then divide by its length and then average over LCPs. Each of these possibilities may give a slightly different result. We choose to work with total bonus over total cost since it is a simple and unambiguous metric to both define and measure.

Definition 4 The average path bonus in a given graph G is the total bonus divided by the number of paths, $\frac{TB}{\# \text{ paths}}$, similarly the average path payment is $\frac{TP}{\# \text{ paths}}$ and the average path cost or the average distance is $\frac{TC}{\# \text{ paths}}$.

Note that the average path bonus can readily be inferred from the average overpayment TB/TC through multiplying the latter by the average distance, $TC/\# \text{ paths}$.

We need to be careful when extending the overpayment definition to a family of random graphs.

Definition 5 The average VCG overpayment in a random graph is

$$\text{avg overpayment} = \frac{\mathbf{E}[TP - TC]}{\mathbf{E}[TC]}.$$

We choose this definition instead of $\mathbf{E}[\frac{TP-TC}{TC}]$ for analytical convenience, since in the latter case the expected overpayment may be artificially blown off by a few unlikely graphs with very low total cost. In the case of Erdős-Renyi graphs with a large number of vertices n , TP and TC are well concentrated around their means so both measures give almost the same results in practice.

One of the criticisms of the VCG mechanism is that it may give payments that are much higher than true costs. In the case of the VCG mechanism, this is easy to see in the following example. Imagine a network with only two parallel paths, one consisting of n edges of cost 1 and the other

consisting of $2n$ edges of cost 1. In this case, the first path is of lowest total cost, n , however each edge receives a bonus of n so the extra payment to the path becomes n^2 , a factor of n times larger than the true cost. Thus, the average overpayment to this single path is essentially as large as the size of the graph.

This example suggests that even short paths may receive a very high overpayment, if the second best paths are much longer than the LCP. This may altogether defeat the practical use of the VCG mechanism. A recent calculation [11] shows, however, that the average overpayment in the Internet graph is only 0.44 (assuming all edge costs are the same), which is incredibly low in light of the size of the Internet and the above pessimistic example. Feigenbaum *et al.* [11] explain this low overpayment by relating it to competition in the presence of network formation: as soon as a link gets a high price, a competing provider will establish another link to capture some of the revenue, thereby reducing overall payments.

As noted by Elkind [9], the economics literature typically requires that each payment to an edge is in the support of its cost distribution. Thus if edge costs are bounded, payment should not be higher than the upper bound on the costs. However, this also requires that the mechanism knows the support or cost distribution of each agent, something that will not necessarily hold in reality. Thus, we will not cap payments and will preserve the payment structure specified by Eq. 1 regardless of cost distributions.

Theoretical insight is often gained through the simplest examples, and in light of this we study the case of unit costs. We also study a scenario of complete graphs with uniform edge costs in [0, 1].

3 VCG Overpayment in the Graph $G(n, p)$

3.1 A formula for VCG overpayment

We begin by studying VCG overpayment in the Erdős-Renyi random graphs with n nodes and edge probability p , $G(n, p)$. Their structure and equal node degree $(n - 1)p$ have been considered unrealistic, in view of the recent literature on the power-law structure of small world networks, the collaboration graphs, the Internet graph, etc. [6, 14]. Nevertheless, the Erdős-Renyi graphs provide important insight for studying more general random graph models, and were used until recently to generate various network topologies [10].

In the entire section, we restrict attention to the case of unit edge costs.

Recall that $LCP(G, u, v)$ stands for (the cost of) the lowest-cost path between u and v and whenever the graph in question is clear from the context, we will simply write $LCP(u, v)$. Also recall that in a random graph, the average overpayment is the ratio of expected total bonus and total cost, $\mathbf{E}[TB] / \mathbf{E}[TC]$.

Theorem 2 *For $G \in G(n, p)$ with $np = \omega(\sqrt{n \log n})$, with probability $\Omega(1 - n^{-c})$ for some constant $c > 0$, the average VCG overpayment is $\frac{p}{2-p}$.*

Proof: Our proof is based on the following two observations.

Claim 3 *For $G \in G(n, p)$ with $np = \omega(\sqrt{n \log n})$, with probability $\Omega(1 - n^{-c})$ for some constant $c > 0$, the graph G has at least two length 2 paths between every pair of vertices.*

Claim 4 *For any graph with two length-2 paths between every pair of vertices, the total bonus is m and the total cost is $n(n - 1) - m$, where n and m are the number of nodes and edges in the graph.*

From here, the result follows immediately by noting that the expected number of edges in $G(n, p)$ is $n(n-1)p/2$, so the average overpayment in a random Erdős-Renyi graph is

$$\frac{\mathbf{E}[m]}{\mathbf{E}[n(n-1) - m]} = \frac{n(n-1)p/2}{n(n-1) - n(n-1)p/2} = \frac{p}{2-p}.$$

We first prove Claim 4. Assume we have a graph with two length 2 paths between every pair of vertices. Then for every pair of vertices $\{u, v\}$, there are two possibilities:

- There is an edge (u, v) . Therefore $LCP(u, v) = 1$, $bonus(u, v) = 1$ since the second shortest path between u, v is of length 2, so the contribution of (u, v) to TB is 1 and to TC is 1.
- There is no edge (u, v) . Therefore $LCP(u, v) = 2$, $path\ bonus(u, v) = 0$ since the second shortest path between u, v is also of length 2. So the contribution of (u, v) to TB is 0 and to TC is 2.

Therefore, if the number of edges in the graph is m , the total cost is

$$TC = 1 * m + 2 * \left(\frac{n(n-1)}{2} - m \right) = n(n-1) - m$$

and the total bonus is $TB = m$. This concludes the proof of Claim 4. Note that the average overpayment in this particular graph is given precisely by

$$\frac{TB}{TC} = \frac{m}{n(n-1) - m}.$$

We now turn to Claim 3. We will show that when $p^2 > k_c \log n/n$ for a constant $k_c = k > 0$ to be specified later, any pair of vertices in a random graph $G(n, p)$ are connected by two length 2 paths with probability $\Omega(1 - n^{-c})$ for a given $c > 0$.

Consider vertices u, v . The probability that they are both connected to a given third vertex w is p^2 and so the probability that they are connected to exactly k of the remaining $n-2$ vertices is

$$\binom{n-2}{k} (p^2)^k (1-p^2)^{n-2-k}.$$

In particular, the probability that u and v are connected by at least two paths of length 2 is

$$1 - (1-p^2)^{n-2} - (n-2)p^2(1-p^2)^{n-3},$$

which approaches 1 as n grows to infinity and p is constant.

Now suppose $p^2 = k \frac{\log n}{n}$. Then the probability that there are less than two length 2 paths between a given pair of nodes is

$$\begin{aligned} (1-p^2)^{n-2} + (n-2)p^2(1-p^2)^{n-3} &= \left(1 - \frac{k \log n}{n}\right)^{n-2} + (n-2) \frac{k \log n}{n} \left(1 - \frac{k \log n}{n}\right)^{n-3} \\ &\leq e^{k \log n} + k(\log n) e^{k \log n} \\ &= (1 + k \log n) \frac{1}{n^k} \leq \frac{1}{n^{k-2}}. \end{aligned} \tag{2}$$




Figure 1: Average overpayment in $G(n, p)$.

p	Edges	TC	TC/path	TP	TP/path	TB	TB/TC
0.1	496	11011	2.22	15299	3.1	4288	0.39
0.2	1001	8962	1.81	10495	2.1	1533	0.17
0.3	1493	8407	1.70	9910	2.0	1503	0.18
0.4	2002	7898	1.60	9900	2.0	2002	0.25
0.5	2471	7429	1.50	9900	2.0	2471	0.33
0.6	3002	6898	1.39	9900	2.0	3002	0.44
0.7	3459	6441	1.30	9900	2.0	3459	0.54
0.8	3938	5962	1.20	9900	2.0	3938	0.66
0.9	4464	5436	1.10	9900	2.0	4464	0.82
1.0	4950	4950	1.00	9900	2.0	4950	1.00

Table 1: Average VCG overpayment for $G(n, p)$ with 100 nodes

Since the probability above is a decreasing function in p , the same bound would apply for $p^2 > k \frac{\log n}{n}$. By the union bound, the probability that any of the pairs in the graph G have less than two length 2 paths connecting them is upper-bounded by

$$\binom{n}{2} \text{Prob}(\text{given pair has less than two length 2 paths}) \leq \frac{n^2}{2n^{k-2}} \leq \frac{1}{n^{k-4}}. \quad (3)$$

Setting $k = c + 4$ for the given constant c concludes the proof of the claim and of our theorem. \square

Note that a higher threshold for p would improve the probability of occurrence of two length 2 paths between every pair of vertices, on the other hand our simulation results (see Figure 1) suggest that the threshold for which the formula starts to hold occurs approximately at $p = \sqrt{\frac{\log n}{n}}$.

3.2 Average VCG payments in $G(n, p)$

Note that average VCG overpayment, defined as the ratio of total bonus to total cost, is also equal to the ratio of the average path bonus to average path cost (dividing the numerator and

denominator by the number of vertex pairs in the graph). From Theorem 2, when $p > \sqrt{\log n/n}$,

$$\frac{TB}{TC} = \frac{\text{Avg Path Bonus}}{\text{Avg Path Cost}} = \frac{p}{2-p},$$

with high probability. When we are assured of having at least two paths of length 2 between any given pair, the average path bonus is precisely p and the average path cost is $2-p$. This is because the average path cost is 1 with probability p , when there is an edge between the pair of nodes, and 2 with probability $1-p$. Similarly the average path bonus is 1 with probability p and 0 with probability $1-p$.

Consequently, the average path payment, which is the sum of the average path bonus and average path cost, is always $p + (2-p) = 2$.

Corollary 1 *With high probability, when $p > \sqrt{\log n/n}$, the average path bonus is p and the average path payment is 2.*

Lastly, the average payment per unit cost is exactly 1 plus the average overpayment, namely $\frac{2}{2-p}$. So just like the average overpayment, average payment also rises when p does. We are now armed to explain this counterintuitive phenomenon. Whenever p is above the given threshold which guarantees the existence of paths of length 2, the bonus of each edge is upper-bounded by 1. On the other hand, in terms of overpayment, it is better to have two paths of length 2 than an edge and a second path of length 2 between a pair of vertices, since in the first case even though the lowest cost path is longer, it gets zero bonus. Similarly, the average payment per unit cost is higher because of the presence of more edges, and not because of having longer LCPs of length 2.

To appreciate these theoretical predictions, we have included Table 1 with simulation results for random graphs with $n = 100$ nodes. We see that initially the total payment TP and the average path payment, TP/path, fall (with the increase in competition for small values of p) and afterwards as predicted, they remain constant. Also the TC/path column which shows the average distance, becomes equal to $2-p$ for $p \geq 0.2$. For the same values of p , the average overpayment TB/TC, is equal to $\frac{p}{2-p}$. Note that each line in the table corresponds to a single generated random graph with 100 vertices and thus it is remarkable how well concentrated around the mean the results are. In Figure 1, we have plotted the average overpayment for $n = 100$, as well as for several more values of n and see that the greater the number of vertices, the greater the range of p for which the overpayment is equal to $p/(2-p)$. Also note that $p/(2-p)$ is a lower bound on the overpayment, so together the constant upper bound from Mihail *et al.* [12], the average overpayment in $G(n, p)$ is $\Theta(1)$.

3.3 An algorithm for optimal VCG overpayment in $G(n, p)$

Elkind showed recently [9] that it is possible to lower the average VCG payments by deleting a subset of edges from the original network, however it is NP-hard to determine what is the best subset of edges to delete, or even whether a given graph can benefit from edge deletion.

We note in passing that the phenomenon of reducing the VCG overpayment through deleting edges, that is, effectively through eliminating competition, is similar to Braess's Paradox for traffic flow. The latter states that sometimes closing down roads may help lower traffic delays [5].

While in general it is hard to determine the optimal set of edges to delete, in the case of $G(n, p)$ our results on overpayment and the simulation results in Fig. 1 suggest a simple strategy to do so.




Figure 2: Average VCG overpayment vs $(\log N)^\alpha$ in complete graphs with uniformly random edge costs. Linear regression fits: (left) $-.291 + .341(\log N)^{1.5}$; (right) $-.118 + .902(\log N)^2$;

First observe from Figure 1 that the lowest VCG overpayment occurs at a threshold value for p , approximately when overpayment becomes $\frac{p}{2-p}$. This threshold occurs approximately at $np = \sqrt{n \log n}$. We offer the following simple randomized algorithm. If the current average degree is below the threshold value, add edges uniformly at random until the average degree reaches the threshold and vice versa, if the current average degree is above threshold, delete edges uniformly at random, until we reduce the degree to that level.

4 Complete graph with random edge costs

Once we have studied and understood well the overpayments in a random Erdős-Renyi graph with unit costs, it is natural to extend this to non-unit costs. We may ask for example, what happens when costs are uniform on $[0, 1]$.

While we can simulate the Erdős-Renyi graph model in this case and make further observations, we will instead look at a seemingly easier case, the complete graph with uniform costs. Cooper *et al.* [7] showed that with high probability, shortest paths in this model consist of $O((\log n)^2)$ edges and the diameter is $O(\log n/n)$.

Note that with uniform costs it is no longer the case that payment per edge and payment per unit cost are the same.

The complete graph is in a sense perfectly competitive since it is completely connected. Thus, we may intuitively expect a low overpayment. However, just like intuition belied us before, this case turns out to be worse than even the sparse Erdős-Renyi graph with unit costs, which has overpayment less than 1.

We simulate the overpayment in complete graphs with $n = 10$ to 450 nodes. Our results are that the average overpayment is strictly increasing with n . The complete graph with only 10 vertices has overpayment of 1.46 and the graph with 450 nodes has overpayment of 5.16, *i.e.*, over 500% of true cost. We study the rate of increase and plot the regression line fits of overpayment vs $(\log n)^{1.5}$ and $(\log n)^2$ in Figure 2. We also run a linear regression of the overpayment versus $(\log n)^{2.5}$ and $(\log n)^3$ and get regression coefficients $.894 + .046(\log N)^{2.5}$ and $1.214 + .018(\log N)^3$ respectively.




Figure 3: Power-law graphs with unit costs. Left: Average VCG overpayment, lower-bounded by $\Theta(\log \log n)$. Right: Average path bonus seems to grow more slowly than $\Theta(\log n / \log \log n)$.

It is hard to estimate the precise asymptotical behavior of the overpayment just from the generated graphs with up to 450 vertices, although based on these regression results $\Theta((\log n)^2)$ looks like a good candidate. In any case, the overpayment seems to be lower-bounded by $\Omega((\log n)^{1.5})$

The average path cost falls a little faster than the diameter bound $O(\log n/n)$ from Cooper *et al.* [7]. For $n = 10$, the average path cost is 0.252 and it falls monotonically down to 0.015 as the number of vertices increases to $n = 450$. The average path bonus also falls monotonically from 0.329 for $n = 10$ to 0.075 for $n = 450$, although it falls at a slower rate than the average path cost.

5 Power-law Graphs

Perhaps the most common real world networks are scale-free, that is their degree sequence follows a power-law distribution: $\Pr(\text{degree} = d) \propto 1/d^\beta$, where typically $2 < \beta < 3$. The Internet graphs have β between 2.1 and 2.45 [14], while the collaboration graph of *Mathematical Reviews* has $\beta = 2.97$ [6].

Chung and Lu have pioneered the study of random graphs with given expected degrees and have shown that the average distance of a power law random graph with exponent $2 < \beta < 3$ is $O(\log \log n)$. The average distance changes to $\Theta(\log n / \log \log n)$ when $\beta = 3$ and to $O(\log n)$ when $\beta > 3$ [6]. Mihail, Papadimitriou and Saberi concluded that as an immediate corollary of Chung and Lu's result, the average VCG overpayment is bounded by $O(\log \log n)$ when $\beta < 3$ [12]. They also state the conjecture that the overpayment in this case is constant.

The other popular model for generating scale-free graphs is the preferential attachment model, which generates graphs by connecting each new vertex to the already existing vertices with probabilities proportional to their current degrees. Bollobas and Riordan [3] offer a precise definition of the model, which was first suggested by Barabasi [2].

The standard preferential attachment model has a parameter M , which stands for the number of vertices that each newly arriving vertex attaches to. This parameter determines the average degree. Since each new vertex adds M new edges, the total number of edges in the graph is approximately Mn , where n is the number of nodes. If the average degree is d , the total number of edges is $nd/2$

since each edge is counted twice. Thus, $d \approx 2M$.

The average degree in the Internet AS-level graph (essentially, the graph of Internet Service Providers) from September 2003 is 4.12, its maximum degree is 2415 and its power-law exponent is about 2.17 [13]. Based on the average degree it is most appropriate to simulate the graph with the preferential attachment model with $M = 2$. Note however, that it is impossible to match all parameters such as average degree, power-law exponent of the degree sequence, maximum degree, etc. Bollobas *et al.* have shown that regardless of M , the preferential attachment model generates a degree sequence with a power-law distribution with exponent 3 [4]. Despite these differences from the real-world network properties, we base our simulations on the preferential attachment model with $M = 2$ since it is easy to generate graphs and also to obtain their biconnected components simply by pruning the vertices of degree 0 and 1. In our results, we state the original number of nodes n that we simulate, however the overpayment and other metrics are based on the biconnected component which contains about 80% of the original vertices.

5.1 Power-law with Unit Costs

Bollobas and Riordan [3] proved that power-law graphs generated by the preferential attachment model with $M \geq 2$ have diameter $O(\log n / \log \log n)$. This bounds the average distance as well. Recall, these graphs all have exponent $\beta = 3$ of the power-law degree sequence. On the other hand, Chung and Lu [6] have shown that in the expected degree model, power-law graphs with $\beta = 3$ have diameter almost surely $\Theta(\log n)$ and average distance $\Theta(\log n / \log \log n)$. Our simulation results for the average distance are consistent with both models, as they appear to grow as $\Theta(\log n / \log \log n)$.

It is remarkable that even though the average distance is higher in this case ($\beta = 3$) than the case $2 < \beta < 3$, in which the distance is $O(\log \log n)$ in the expected degree model [6], the average VCG overpayment is still not only constant but steadily declining. In Figure 3 we plot the overpayment, as well as $\frac{29}{30} * 1 / \log \log n$ for a comparison; we see that $\log \log n$ seems to decrease more rapidly than the overpayment, so the overpayment is bounded below by $\Omega(1 / \log \log n)$. In addition, the overpayment is bounded above by 0.5 for all n in our simulation. We also obtain that the average path bonus is increasing with n . For example, when $n = 100$, the average path cost is 2.59 and the path bonus is 1.28. They increase to 4.56 and 1.95 respectively when $n = 20,000$. Naturally, the average distance increases with n and as paths get longer, they receive higher payments, and as it turns out higher bonuses.

Recall that the average overpayment and the average path bonus are related as

$$\frac{TB}{TC} = \frac{\text{avg path bonus}}{\text{avg path cost}},$$

that is, average overpayment is equal to the average path bonus divided by the average distance. Since the average distance is $O(\log n / \log \log n)$ and the average overpayment TB/TC is $O(1)$, then the average path bonus is also $O(\log n / \log \log n)$. This is just an upper bound, and is consistent with the plot on the right in Figure 3. If the average distance is exactly $\Theta(\log n / \log \log n)$, then by our estimates above the average overpayment is asymptotically lower bounded by $1 / \log \log n$ and so the average path bonus is $\Omega(\log n / (\log \log n)^2)$.

5.2 Power-law with Uniform Random Costs

Recall that in complete graphs, overpayment is exactly 1 with unit costs and exceeds $\Theta((\log n)^{1.5})$ with uniformly random costs, thus the cost distribution alone can exacerbate payments significantly.

Somewhat surprisingly, we obtain that the average overpayment in our simulated power-law graphs with uniformly random costs remains 0.85 ± 0.01 as n increases from 100 to 4,000 so it seems to be constant for any n .

These results suggest that the average VCG overpayment is bounded by constant in power-law graphs, regardless of the value of β , and regardless even of whether edge costs are unit or uniformly random. It is an intriguing open question to prove this theoretically, for either the preferential attachment or expected degree model, and general cost distributions.

6 Concluding Remarks

In this paper, we studied theoretically and empirically the VCG overpayment in different types of graphs with unit and uniformly random edge costs. We started out by defining two main notions of overpayment, average VCG overpayment per unit cost (which we call simply average overpayment) and average path bonus.

We derived a formula for average overpayment in the classical Erdős-Renyi random graphs. This formula helped explain a counter-intuitive phenomenon that overpayment increases when graphs become denser and hence more competitive. We carried through this observation in complete graphs with uniformly random edge costs, which have the worst overpayment of all graphs we consider, of more than $\Omega((\log n)^{1.5})$. Our empirical results on power-law graphs showed that the overpayment decreases with the size of graphs with unit costs and stays constant in graphs with uniformly random costs.

In conclusion, our results suggest that for common real-world networks, the VCG payments will not be too much higher than the true costs, hence the VCG mechanism should remain an important and attractive option for price or incentive integration to the Internet and other networks.

Acknowledgment

We thank Rahul Sami, Edith Elkind and Amin Saberi for stimulating discussions. We also thank Yifei Chen for pointing out a missing detail in the original proof of Theorem 2.

References

- [1] A. Archer and E. Tárdoš. Frugal Path Mechanisms. SODA 2002, pp.991 – 999.
- [2] A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science* 286(1999), 509 – 512.
- [3] B. Bollobas and O. Riordan. The diameter of scale-free graphs. *Combinatorica*, To appear.
- [4] B. Bollobas, O. Riordan, J. Spencer, and G. Tusnady. The degree sequence of a scale-free random graph process. *Random Structures and Algorithms*, 18(3):279–290, 2001.
- [5] D. Braess. Uber ein paradoxon der verkehrsplanung. *Unternehmensforschung*, 12:258-268, 1968.
- [6] F. Chung and L. Lu. The average distance in a random graph with given expected degree sequences. Proceedings of the National Academy of Sciences, 99, No 25, December 2002.
- [7] C. Cooper, A. Frieze, K. Mehlhorn, V. Priebe. Average-case complexity of shortest-paths problems in the vertex-potential model. Randomization and approximation techniques in Computer Science (*Proceedings of RANDOM '97*). Lecture Notes in Computer Science 1269, 15-26.
- [8] E. Elkind, A. Sahai and K. Steiglitz. Frugality in Path Auctions. *Proc. 15th ACM-SIAM Symposium on Discrete Algorithms*, 2004.
- [9] E. Elkind. True Costs of Cheap Labor are Hard to Measure: Edge Deletion and VCG Payments in Graphs. Manuscript, April 2004.

- [10] L. Li, D. Alderson, W. Willinger and J. C. Doyle. A First-Principles Approach to Understanding the Internet's Router-level Topology. *Proceedings of ACM Sigcomm*, Portland, OR, August 2004.
- [11] J. Feigenbaum, C. Papadimitriou, R. Sami and S. Shenker. A BGP-based Mechanism for Lowest-Cost Routing. *Proceedings of the 21st Symposium on Principles of Distributed Computing*, ACM Press, New York, 2002, pp. 173 – 182.
- [12] M. Mihail, C. H. Papadimitriou, A. Saberi. On Certain Connectivity Properties of the Internet Topology. FOCS 2003.
- [13] Routeviews routing table, snapshot from Sep.15, 2003, 18:00 h. <http://www.routeviews.org/>
- [14] G. Siganos, M. Faloutsos, P. Faloutsos and C. Faloutsos. Power-Laws and the AS-level Internet Topology. *Transactions on Networking*, to appear.
- [15] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker and W. Willinger. Network Topology Generators: Degree-Based vs. Structural. SIGCOMM 2002.