Boycotting and Extorting Nodes in an Internetwork

Dave Levin Adam Bender Cristian Lumezanu Neil Spring Bobby Bhattacharjee

University of Maryland, College Park
Emails: {dml,bender,lume,nspring,bobby}@cs.umd.edu

1. INTRODUCTION

A *boycott* is a protest or a demand for change based on consumers following a simple rule: do not purchase service or goods from specific producers. In any network, and especially in the Internet, such grass-roots protest is close to impossible. The big networks do not connect directly to consumers and routing is based on locally choosing the shortest, cheapest, or most profitable paths. In this environment, individual consumers have no influence on route selection and enterprise customers very little.

Yet, protest-worthy actions by network providers are not infrequent. AOL compromised the privacy of its users by publishing search queries (August 2006); Google complied with China's censorship laws, to the disapproval of many (February 2006); Verisign redirected mistyped DNS names to their own advertising (September 2003); and the Electronic Frontier Foundation has sued AT&T to stop NSA surveillance (January 2006; legal action continues).

We envision an Internet in which users have the ability to make, or at least influence, routing decisions. Users will have this ability so that they may take advantage of increasingly rich functionality in networks; users will need to be able to choose networks that provide a service or avoid those that filter or censor. Among the many policies users might choose in routing packets, users may select nodes to boycott.

To study the economics of boycotting a provider in an abstract network, we adapt the Vickrey-Clarke-Groves (VCG) routing mechanism to support boycotting. All link costs are globally advertised. Sources in this network typically choose the shortest cost path to a destination, and micropayments reward each entity along the path. Each of these assumptions is distinct from the realities of BGP; we describe the implications of this mismatch late in this paper.

We begin the investigation into the following questions: (1) Does VCG continue to encourage nodes to report costs truthfully in the presence of boycotting? (2) Do nodes gain from learning users' boycott lists? Do users gain from divulging whom they are boycotting? and (3) How do the results from VCG apply to ad hoc networks and the Internet?

To answer these questions, we modify VCG (§2) and show that only when massively boycotted do nodes have incentive to lie about link costs (§3). We simulate our modified VCG (§4) to show the price in connectivity, and discuss deployment requirements and implications (§5). We review related work (§6) and conclude (§7).

2. A MECHANISM FOR BOYCOTTING

2.1 Model, Assumptions, and Problem

We model the network topology as a directed, weighted graph $G = (V, E)$. For any two nodes $u, v \in V$, $c(u, v)$ is the nonnegative cost incurred by u to send on the directed edge $(u, v) \in E$. A node can have at least one of the following roles: source, destination, or transit. Each source node s has a (possibly empty) set $\mathcal{B}_s$ of nodes that it wishes to boycott. That is, s prefers that no traffic sent by s follows a path that contains any $b \in \mathcal{B}_s$.

We assume that source routing is supported, but show next that *source routing alone is not enough to boycott.* We assume a *payment mechanism* by which s can securely pay other nodes in the network, and an *accountability mechanism* with which s can verify that its source route was used. For accountability, each node might be required to sign the packets it forwards, but this strawman approach verifies only that the packet visited *at least* those nodes that signed it. In this section, the central challenge is to allow a source to compel the nodes on its chosen path to forward its packets along only that chosen path, even when more profitable alternate paths exist.

2.2 VCG Routing Mechanism

We first review the Vickrey-Clarke-Groves (VCG) mechanism [6, 9, 19]. Each transit node i knows the cost $c(i, j)$ to send to node j where $(i, j) \in E$. A source s routes to a destination d by first requesting from each transit node its costs, $c(i, \star)$. In response, i claims cost $c'(i, j)$ for all $(i, j) \in E$. The goal of VCG is to ensure that every $c(i, j) = c'(i, j)$: that each i reports its costs *truthfully.* Let $\mathrm{SCP}(s, d, H)$ denote the shortest cost path from s to d in graph H, $\mathcal{C}(\mathrm{SCP}(s, d, H))$ the sum of the edge costs on the path, and let H_{-i} denote the graph induced by removing i from H. VCG accomplishes truthful reporting of costs by employing the following payout function from s to i: if i is not on the shortest cost path, and thus does no work, $p_s(i) = 0$, otherwise:

$$p_s(i) = \mathcal{C}(\mathrm{SCP}(s, d, G_{-i})) - \mathcal{C}(\mathrm{SCP}(s, d, G)) + c'(i, j) \quad (1)$$

We summarize VCG with the pseudocode in Algorithm 1. To see why this induces truthful reporting on behalf of each transit node i, observe that i should neither understate nor overstate its cost:

i should not understate: By definition, $\mathcal{C}(\mathrm{SCP}(s, d, G_{-i})) \geq \mathcal{C}(\mathrm{SCP}(s, d, G))$, so each node on the shortest cost path receives at least its *stated* cost to forward. If $c'(i, j) <$

$c(i,j)$, then i is not guaranteed to recover its actual cost to forward, hence i should ensure that $c'(i,j) \geq c(i,j)$.

i **should not overstate**: $p_s(i)$ is independent of $c'(i,j)$ because $\mathcal{C}(\mathrm{SCP}(s,d,G))$, in Eq.(1), includes $c'(i,j)$. Increasing $c'(i,j) > c(i,j)$ could, however, remove i from the shortest path, so i should ensure that $c'(i,j) \leq c(i,j)$.

Algorithm 1 Standard VCG

1. Gather $c'(i,j)$ from each i for each $(i,j) \in E$.
2. Compute the shortest cost path P from s to d using Dijkstra's algorithm.
3. For each intermediate hop i on P, pay $p_s(i)$ from Eq.(1) to i.

2.3 VCG with Boycotting Mechanism

Culled-VCG. Suppose source s wishes to boycott a set of nodes $\mathcal{B}_s$. One could run VCG on the graph induced by removing all boycotted nodes, $G_{-\mathcal{B}_s}$; we refer to this as *culled-VCG*. This approach is similar to MIRO [20], which employs source-routing to allow for nodes to visit (or avoid) nodes of their choosing. We show in Figure 1, however, that culling the boycotted node alone does not suffice. In it, node s is sending to node d but wishes to boycott B. The shortest cost path computed by standard VCG is $\{s,A,B,C,d\}$, and by culled-VCG, $\{s,A,C,d\}$. A, however, could gain more profit by tunneling s's packets through B at a cost of $2+\epsilon$, giving B profit ϵ, and giving A profit $(1-\epsilon)$ in addition to the profit made from s's payment. Culled-VCG fails because, *although nodes have incentive to truthfully report their costs (thanks to VCG), it provides no incentive for nodes to forward on the path the source specified.*

VCG with Boycotting. We can draw from culled-VCG the following observation: for each edge (i,j) in the shortest cost path computed by culled-VCG, if there is a lower-cost path from i to j that includes a boycotted node (such as the (A,C) edge in Figure 1), then the packets will most likely go through the boycotted node. We introduce a modified VCG mechanism in Algorithm 2.

Algorithm 2 VCG with boycotting

1. Gather $c'(i,j)$ from each i for each $(i,j) \in E$.
2. Compute all-pairs shortest-cost paths on the weighted graph $G=(V,E,C)$, with the weight of edge (i,j) equal to $c'(i,j)$.
3. Let $G'=(V,E',C)$, where $E' = \{(i,j) \in E \mid$ the least-cost *path* from i to j does not contain any $b \in \mathcal{B}_s\}$.
4. Compute the least-cost path P from s to d using Dijkstra's algorithm on $G'_{-\mathcal{B}_s}$ (i.e., run culled-VCG on G').
5. For each intermediate hop i on P, pay $p_s(i)$ from Eq. (1) to i.

2.4 The Cost of Boycotting

Boycotting a node is likely to incur additional cost. When s applies Algorithm 2 to Figure 1, the resulting path is $\{s,X,Y,d\}$, with payouts $p_s(X)=p_s(Y)=4$, and $c(s,X)=2$, for a total cost to s of 10. Were s to apply standard VCG, the path would have been $\{s,A,B,C,d\}$, with payouts $p_s(A)=p_s(C)=3$, $p_s(B)=2$, and $c(s,A)=1$, for a total cost of 9. We define the *cost of boycotting* to be the ratio of the total price paid by s when running VCG with boycotting (Alg. 2) to the price paid when running standard VCG (Alg. 1). In the example of Figure 1, s has a cost of boycotting of $10/9$ when sending to d and boycotting B.

If s values boycotting infinitely, then s is willing to accept any cost of boycotting, even at the cost of not being able to reach the destination. However, if, for each $b \in \mathcal{B}_s$, s assigns some finite utility $u(b)$ to boycotting b, then s may not be willing to boycott if the difference in price is greater than $u(b)$. In such a situation, s may be willing to accept some subset S of $\mathcal{B}_s$ to boycott such that the additional cost incurred to boycott S is less than the utility s gains, $\sum_{b \in S} u(b)$. We leave combinatorial boycotting policies (e.g., "send to no nodes in $S \subset \mathcal{B}_s$, or no more than 3 nodes from $T \subset \mathcal{B}_s$") to future work.

3. BEHAVIOR IN A NETWORK WITH BOYCOTTING

Beyond the route determination game, there are additional strategies that both boycotters and boycotted nodes must make. As long as boycotters do not value boycotting infinitely, boycotted nodes may make rational, competitive responses. First, a transit node i can attempt to determine if it is being boycotted, why, and the impact to i's profits. Second, i may attempt to regain market share, either via policy change (so as to stop being boycotted) or by extorting those who are not boycotted. We address these in turn, and finish this section with a discussion of how nodes are expected to report their costs.

3.1 Knowing Your Boycotters

Routing entity i may wish to learn about users' boycott sets. For instance, i may ask: (1) Are users boycotting me due to a specific policy I do (or do not) exhibit? (2) How many users are boycotting me, and how much profit am I losing as a result?

Boycotted nodes may find it extremely difficult to answer these questions without being able to view the traffic that they would otherwise carry. In a wired network, when i is boycotted, it receives no packets, and hence, without a priori knowledge of traffic distributions, has no way of inferring how much traffic it is not seeing. In omnidirectional wireless networks, such as 802.11, each transmission is effectively also a broadcast, so i could perhaps learn more of the existing traffic flows, but only in close physical proximity. Still, i has several tools at its disposal: the observed amount of traffic and payments

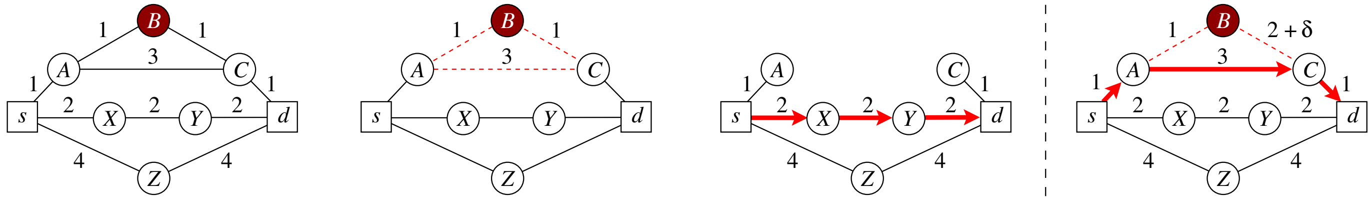

Figure 1: (a) Example network: s wishes to boycott B. (b) Edges that would involve B are culled. (c) The shortest path in the resulting graph is chosen with cost of boycotting $10/9$. (d) An example where B could extort A.

it has received with a given $c'(i,\star)$, and the stated (and hence truthful) costs of others' edges.

Our insight is that i can periodically perform "market research" by lowering its stated costs[1] and compare the traffic it is asked to forward during the lower-cost period to normal periods. Suppose, for instance, i receives packets with source-destination pair (s,d) only when it states a lower cost, but i is on the least-cost path from s to d even when i is truthful. Then i has revealed an instance where s is boycotting i but, upon stating lower cost, i induces a cost of boycotting that is too high for s to cover, at which point it is not economically feasible for s to boycott i.

3.2 Eliciting Change

Why would a boycotted node have to perform this market research? Would a boycotter not want to be outspoken as to whom and why it is boycotting?

Indeed, boycotting is a method to elicit policy change. If enough end-users boycott a transit node, the message should be clear: change policy to earn greater profits. Boycotting is, in a sense, a means of resolving the tussle between end-users who want nodes that forward their packets to enforce a particular policy (or lack thereof) and the routing entities that do or do not support that policy. To do so, boycotters must make clear to the routing entities *why* they are being boycotted.

In two cases, however, a node s may wish to "silently" boycott some routing entity i. First, s may be sending confidential material, using the boycott as a means to avoid snooping or to access diverse paths, and does not want to publicly announce that it is boycotting i to do so. Second, s may wish to help the boycott gather momentum and become significant before stating grievances, hoping to deny i the ability to quell the boycott with half-measures. Hence, market research as we have described it may serve powerful routing entities who wish to learn about growing, grass-roots boycotting efforts.

3.3 Extorting Non-Boycotted Nodes

As an alternative to changing its policy, routing entity i may extort other, non-boycotted routing entities to gain profit. Consider the example in Fig. 1(d); node B's true cost to send to C is $2+\delta$ for some positive δ. In this case, our mechanism (Alg. 2) would not remove edge (A,C), and the path $\{s,A,C,d\}$ would be chosen, with payment to A, $p_s(A) = (6-5+3) = 4$. A would receive a profit of 1 (after paying its cost of 3 to forward). B may threaten to *understate* its cost to 1 (resulting in Fig. 1) unless A forwards some amount $\epsilon < 1$ of its profit to B.

A has incentive to pay B (and earn $1-\epsilon$) rather than allow B to understate its cost (and earn 0). Of course, by understating its cost, B runs the risk of not covering its cost to forward traffic for other (s,d) pairs for which B **is** on the least-cost path. If the total profit from extorting is greater than the loss in profit from understating, then B has incentive to extort.

Source node s may wish to prevent $b \in \mathcal{B}_s$ from extorting for the same reasons that it wishes to boycott b: to protect against b earning money or to make it economically infeasible for b to inspect s's packets. To prevent against extortion attacks, we modify Algorithm 2 as follows:

Algorithm 3 VCG with boycotting and anti-extortion

1. Gather $c'(i,j)$ from each i for each $(i,j) \in E$. If $i \in \mathcal{B}_s$, ignore the stated cost and set it to 0 instead.
2. Run steps 2–5 of Algorithm 2 as stated.

The intuition behind this modification is that s ought to provision for the worst-case scenario, in which b threatens to understate its costs as much as possible (to zero). If, in so doing, b would not modify the chosen path, then b cannot extort any nodes on the path.

3.4 Cost Revelation

We investigate whether Alg. 3 gives incentive to node i to truthfully state its edge costs.

3.4.1 *Boycotted Nodes*

Suppose $i \in \mathcal{B}_s$. As discussed above, i could employ policy change to try to remove itself from $\mathcal{B}_s$, but can i gain any utility from modifying its stated cost? Even if i were to know that he is being boycotted by s, the answer to this question is "no" by construction; the first step s takes is to ignore i's stated costs. This removes even the possibility of extorting other nodes, as i has no leverage over s's routing decisions. Thus, with respect to any node s with $\mathcal{B}_s \ni i$, i's strategy set is effectively nil, and truth-telling is a weakly dominant strategy (there are no other strategies that improve i's payout).

[1]perhaps for only a short period of time, so as not to incur too great a loss

3.4.2 Non-Boycotted nodes

Our mechanism is powerful in that it is impervious to any countermeasures made by the boycotted nodes. We show here that, surprisingly, it is the nodes that are *not* boycotted who may have incentive to lie about their costs. Fix now an $i \notin \mathcal{B}_s$, and let us consider understating and overstating $c(i,j)$ separately.

Understating Costs. As in standard VCG, for any $i \notin \mathcal{B}_s$, i should not understate its costs. Consider the outcomes of setting $c'(i,j) < c(i,j)$. Since our mechanism uses the same payouts as VCG, we have the same property outlined in Section 2.2; i cannot control the payment that it receives: only whether or not it is on the shortest cost path from s to d, $\mathrm{SCP}(s,d,G)$. If by stating its true cost i would have been on the shortest path from s to d, then i gains no additional utility from understating its cost. If by stating its true cost i was not on $\mathrm{SCP}(s,d,G)$, then stating a $c'(i,j) < c(i,j)$ could put i on the shortest path. However, as described in Section 2.2, the only assurance i has is that $p_s(i) \geq c'(i,j)$, hence i has no guarantee that its cost of forwarding would be recovered, and might therefore decrease its utility.

Overstating Costs. Suppose first that by truthfully reporting its cost, i would be on $\mathrm{SCP}(s,d,G)$. Then by overstating its cost, i would risk not being on the shortest-cost path; further, as above, even if it remained on the path, it would not increase its profit.

In the case where truthful reporting would not place i on $\mathrm{SCP}(s,d,G)$, the proof of standard VCG goes as follows: increasing its cost cannot make its path shorter, and hence i can gain no utility in overstating. However, in our mechanism, there are two cases in which i does not appear in $\mathrm{SCP}(s,d,G)$: In the first case, none of i's edges were culled; i's path simply costs too much to be considered by s, as happened with node Z in Fig. 1. As in standard VCG, such a node experiences no change in utility from overstating. In the second case where, for $i \notin \mathrm{SCP}(s,d,g)$, i would have been on the shortest path but at least one of its edges was culled to avoid a boycotted node, as with node A in the figure. If A were to have overstated its cost to $B \in \mathcal{B}_s$ to some value greater than 2, then s would not have culled the (A,C) edge.

As with the extortion attack, node A will only have incentive to overstate its cost to B if the profit gained from nodes boycotting B is greater than the profit lost from those who would have gone through A to B if the edge cost less. The most likely scenario in which this occurs is when a majority of end-users boycott B.

3.4.3 A Sybil Attack

A transit node can increase the incentive that it has to overstate its costs to boycotted nodes with the following Sybil attack [7]. Suppose node i has n neighbors in graph G. To ensure that i is a viable transit node for any end-user s for which i and at least two of i's neighbors are not in $\mathcal{B}_s$, i can create $\binom{n}{2}$ pseudo-identities ("Sybils"), one for each pair of neighbors. Each such Sybil connects only that pair of neighbors; we show an example in Figure 2.

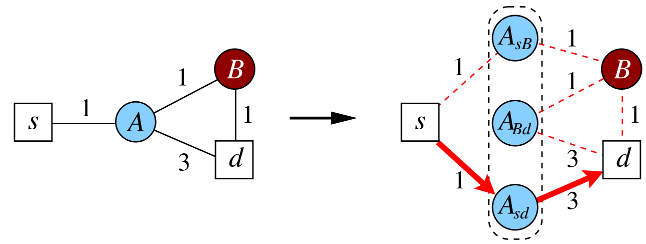

Figure 2: A creates $\binom{3}{2}$ Sybils, at least one of which has an edge to d that s will not cull while boycotting B.

Various schemes protect against Sybil attacks, such as charging each participant an entry fee. We do not attempt to provide a solution here, only to point out that transit nodes may have incentive to perform a Sybil attack and a counter-measure may be necessary to support boycotting. An ISP is unlikely to be able to launch such a Sybil attack.

3.5 Addressing the Overstating of Costs

Unfortunately, our mechanism does not ensure that non-boycotted nodes will truthfully report their edge costs. Although this holds only in extreme cases—when the boycotted node is boycotted by so many end-users that increasing its stated cost would not cause a loss of profit—we would like to have a mechanism that always gives incentive for truthfulness. We believe that there is no payment scheme that will provide this incentive, but are unable to provide a proof at this time. *Addressing overstating is the most important open problem in this line of work.* We are considering heuristics to address this problem, such as: (1) If at any point A is found to be overstating its costs, or if it is found to be forwarding s's packets through some node $b \in \mathcal{B}_s$, then s boycotts A, as well. (2) s could require that A prove its costs to b.

4. SIMULATION RESULTS

To observe the cost of boycotting, we evaluated Algorithms 2 and 3 with a simulator. Our preliminary results indicate that if a node can be boycotted, it can be boycotted cheaply. We evaluate our mechanism both in ad hoc topologies (to test it on general, shortest-cost path networks) and on a subset of the Internet AS topology (which allows only policy-compliant paths).

Wireless Ad Hoc Topology. The simulations used a fixed topology of 130 nodes that was taken from a wireless ad hoc network in Portland, OR [3]. Nodes within 250 meters of each other shared an edge. We assigned edges costs between 1 and 20 from a Zipf distribution (most edges have smaller cost). We chose the Zipf distribution to model a network with many homogeneous links that allow most nodes to communicate cheaply and a small number of more important links whose service comes at a premium. For each ⟨source, destination⟩ pair, chosen at random, we calculated the cost of the shortest path from the source to the destination using VCG. Then, we simulated the source boycotting each node along that shortest path individually, and we ran the boycotting mechanism

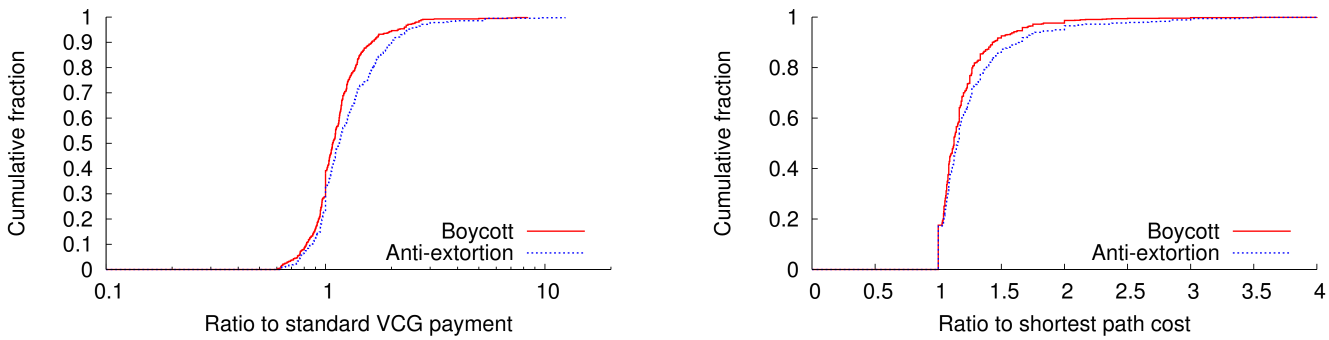

Figure 3: Comparing the boycotting (Alg. 2) and boycotting with anti-extortion (Alg. 3) mechanisms to standard VCG. Left: sum payments made by the source node, normalized to standard VCG. Right: cost of shortest-cost path, normalized to standard VCG. The cost distribution is Zipf.

and the boycotting with anti-extortion mechanism. We present our results in Figure 3.

VCG requires at least two paths between a given source and destination (in order to compute the marginal benefit of each edge on the shortest-cost path), and our proposed mechanisms generally require more (to allow for culled edges). Otherwise, the mechanism will return an infinite cost. In the sample run depicted in Figure 3, VCG returned a finite cost for 62.7% of the ⟨source, destination⟩ pairs, VCG with boycotting returned a finite cost for 51.6% of the ⟨source, destination, boycotted node⟩ tuples (10.6% of the graphs were disconnected), and VCG with boycotting and anti-extortion returned a finite cost for 38.6% of tuples (17.2% of graphs were disconnected).

The plots in Figure 3 are taken across 1082 ⟨source, destination, boycotted node⟩ tuples, and are normalized to the standard VCG costs (when the source is not boycotting). In terms of both VCG cost and path length, the cost of boycotting and the cost of boycotting with anti-extortion are similar. We believe this is because paths are likely to include cheap edges. When a node on the path is boycotted, it is likely that the cost of its outgoing edges is nearly 1, so setting them to 0 (Alg. 3) has little effect on whether or not that path is culled.

For about 25% of paths, the path cost is cheaper when a node is boycotted. This is possible in the following scenario: there is a single cheap path from the source to the destination, and several more expensive paths that have similar costs to each other. When a node on the cheap path is boycotted, a more expensive path must be chosen. Since the expensive paths are priced similarly, the marginal benefit of one over the other is less than the marginal benefit of the cheap path over the expensive paths. In this case, VCG may report that the cost of boycotting is cheaper than when no boycotting is done. This may result in senders boycotting certain nodes to decrease the cost of their paths.

Internet AS Topology. We also ran experiments to try to determine how feasible it was to boycott "major players" in the Internet. For this simulation, we used a graph of the 1170 nodes with highest degree in the Internet AS topology [18]. We set the edge costs of node i to $\lfloor \log_2(\text{degree}(i)) \rfloor + 1$, to model higher-degree nodes as being more expensive, but taking the logarithm to limit the cost disparity. To make sure our results did not depend on the choice of cost function, we also ran separate experiments setting the edge costs of node i to degree (i) and to $\lfloor 1024/\text{degree } (i) \rfloor$. Boycotting many (up to 15) tier-1 ISPs simultaneously, nodes were still able to maintain connectivity for almost all randomly chosen ⟨source, destination⟩ pairs, for all models of cost functions. However, our result almost certainly overestimates the connectivity since we assume that *all* remaining links can be used for forwarding any packet (regardless of origin). In reality, many links will be subject to AS-local policy, and may not be used to forward third-party traffic.

5. IMPLICATIONS IN THE INTERNET

The boycotting mechanism we described is not feasible in today's network. Although there are mechanisms for choosing paths (loose source routing) and for determining which AS-path a next-hop ISP may choose, boycotting is infeasible for three reasons. First, sources do not (and have no mechanism to) pay entities beyond their neighbors. Payment is instead achieved through longstanding, pairwise agreements that are unlikely to change when a boycott begins. Second, the cost of each link is kept secret, and may vary, which means sources cannot reliably estimate the "cost" of preferring one path over another, an estimate that would be needed to compensate the network implementing the route choice. Third, the path a packet takes cannot be verified: a provider might accept payment for boycotting an upstream AS, but proceed to use it anyway.

The effectiveness of a boycott derives from a quantifiable loss in profit that can be traced to a specific event or group of aggrieved sources. Either suddenly shifting traffic away from a boycotted node threatens congestion, or it risks being nullified by traffic engineering.

Overlays, and perhaps loose source routing, provide a mechanism for users to bypass network routing to emulate boycott policies. To boycott many nodes would require many cooperating overlay participants. To boycott with high volumes of traffic may be infeasible, for performance (forwarding may be difficult) and because ISP traffic engineers may readjust traffic away from newly-congested links. We wonder whether an overlay routing approach can (and should) adapt quickly enough to counter any traffic engineering in the middle of the network.

Within BGP, one AS may be able to partially boycott another by inspecting paths advertised by different upstream providers and preferring the provider whose routes do not include the boycotted ISP. The partial boycott does not change BGP in any way. MIRO [20] would provide additional power to this approach by allowing a stub AS to choose among paths offered by its provider. The "cost" of the decision is local, in that the AS changing routes to effect a boycott is no longer optimizing the prior path choice.

Beyond BGP, inter-domain routing proposals have been proposed to allow users or their immediate provider ISPs to exercise greater control over the paths traversed by their packets. Traceback approaches would allow users to verify the path taken by their packets. The combination of source control and verification proposed may be sufficient for implementing any boycotting scheme in a future Internet. The potential flexibility offered by such networks motivates our study of abstract VCG networks.

6. RELATED WORK

Researchers have applied VCG-like [6, 9, 19] mechanisms to improve routing in distributed environments. Nisan and Ronen [14] proved the feasibility of strategyproof, polynomial time VCG mechanisms for networks of selfish users. This result motivated significant work [1, 8, 11, 15]. Anderegg and Eidenbenz propose Ad hoc-VCG [1], a protocol to find minimum-energy paths in mobile ad hoc networks with selfish participants that accept payments for forwarding data. Feigenbaum *et al.* [8] build upon BGP to compute least-cost AS paths. They show that a distributed implementation of VCG does not require considerable cost. These algorithms compute paths that are shortest by a single cost metric and do not consider user-specified criteria for choosing paths.

Providing users with autonomy to choose paths has been proposed as a means to circumvent failures [2] and improve performance [10, 17]. Users or administrators of smaller networks gain this autonomy through source routing [16, 21], overlays [2, 17], or alterations to BGP [13, 20]. These mechanisms may be useful toward implementing a boycott, but the research has been focused on more quantifiable benefits in performance and reliability.

In wireless networks, blacklisting can prevent low quality links or nodes from being considered in the path selection process [4, 5, 12]. However, blacklisting is not done at the request of the sender: a node or link is removed from the network if its performance, by some measure, falls below a predefined threshold.

7. CONCLUSIONS

We introduced the *boycott* as a policy goal applicable to a network. We showed how to adapt the VCG routing model to support boycotting groups of nodes and found that truthful reporting of link costs remains a weakly dominant strategy when boycotting is infrequent. We simulated the routing policies on ad hoc network and AS-graph topologies to show that the cost of boycotting is typically small and that disconnectedness is rare.

Many open problems and research directions remain. Currently, nodes have incentive to overstate their costs to boycotted nodes (effectively disassociating themselves), albeit only when the boycotting is wide-spread. Our work assumes shortest-cost routing, but BGP requires paths to (also) be policy-compliant. Extending our mechanism to work in this setting is crucial to understanding its feasibility in the Internet. Also, we believe boycotting may be used as a means of simulating and estimating the network's resilience to catastrophic failure.

8. ACKNOWLEDGMENTS

This work was partially supported by NSF Awards CNS-626636, CAREER ANI-0092806, and a fellowship from the Sloan Foundation. Dave Levin was also supported in part by a UMD CSD Deans fellowship.

9. REFERENCES

[1] L. Anderegg and S. Eidenbenz. Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents. In *Proc. of MobiCom*, 2003.

[2] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient overlay networks. In *ACM SOSP*, 2001.

[3] H. Balakrishnan, *et al.* The Distance-2 Matching Problem and its Relationship to the MAC-Layer Capacity of Ad Hoc Wireless Networks. *IEEE JSAC*, 22(6):1069–1079, 2004.

[4] M. Busse, T. Haenselmann, and W. Effelsberg. Energy-efficient forwarding schemes for wireless sensor networks. In *WoWMoM*, 2006.

[5] K.-W. Chin, J. Judge, A. Williams, and R. Kermode. Implementation experience with MANET routing protocols. *SIGCOMM CCR*, 32(5):49–59, 2002.

[6] E. Clarke. Multipart pricing of public goods. *Public choice*, 11, 1971.

[7] J. Douceur. The Sybil Attack. In *IPTPS*, 2002.

[8] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In *PODC*, 2002.

[9] T. Groves. Incentives in teams. *Econometrica*, 41, 1973.

[10] K. Gummadi, *et al.* Improving the reliability of Internet paths with one-hop source routing. In *OSDI*, 2004.

[11] J. Hershberger and S. Suri. Vickrey prices and shortest paths: What is an edge worth? In *FOCS*, 2001.

[12] H. Lundgren, E. Nordstro, and C. Tschudin. Coping with communication gray zones in IEEE 802.11b based ad-hoc networks. In *WoWMoM*, 2002.

[13] R. Mahajan, D. Wetherall, and T. Anderson. Negotiation-based routing between neighboring ISPs. In *NSDI*, 2005.

[14] N. Nisan and A. Ronen. Algorithmic mechanism design. In *ACM Symposium on Theory of Computing*, 1999.

[15] D. C. Parkes and J. Shneidman. Distributed implementations of Vickrey-Clarke-Groves mechanisms. In *ACM AAMAS*, 2004.

[16] B. Raghavan and A. C. Snoeren. A system for authenticated policy compliant routing. In *SIGCOMM*, 2004.

[17] S. Savage, *et al.* Detour: Informed Internet routing and transport. *IEEE Micro*, 19(1):50–59, 1999.

[18] http://www.caida.org/data/active/as-relationships/.

[19] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16, 1961.

[20] W. Xu and J. Rexford. MIRO: Multi-path interdomain routing. In *ACM SIGCOMM*, 2006.

[21] X. Yang. NIRA: A new Internet routing architecture. In *FDNA*, 2003.