

Hidden-Action in Multi-Hop Routing

Michal Feldman¹

mfeldman@sims.berkeley.edu

John Chuang¹

chuang@sims.berkeley.edu

Ion Stoica²

istoica@cs.berkeley.edu

Scott Shenker²

shenker@icir.org

¹School of Information
Management and Systems
U.C. Berkeley

²Computer Science Division
U.C. Berkeley

ABSTRACT

In multi-hop networks, the actions taken by individual intermediate nodes are typically hidden from the communicating endpoints; all the endpoints can observe is whether or not the end-to-end transmission was successful. Therefore, in the absence of incentives to the contrary, rational (i.e., self-ish) intermediate nodes may choose to forward packets at a low priority or simply not forward packets at all. Using a principal-agent model, we show how the hidden-action problem can be overcome through appropriate design of contracts, in both the direct (the endpoints contract with each individual router) and recursive (each router contracts with the next downstream router) cases. We further demonstrate that per-hop monitoring does not necessarily improve the utility of the principal or the social welfare in the system. In addition, we generalize existing mechanisms that deal with hidden-information to handle scenarios involving both hidden-information and hidden-action.

Categories and Subject Descriptors

C.2.4 [Computer-Communication Networks]: Distributed Systems; J.4 [Social And Behavioral Sciences]: Economics

General Terms

Design, Economics

Keywords

Contracts, hidden-action, incentives, mechanism design, moral-hazard, multi-hop, principal-agent model, routing

1. INTRODUCTION

Endpoints wishing to communicate over a multi-hop network rely on intermediate nodes to forward packets from the sender to the receiver. In settings where the intermediate nodes are independent agents (such as individual nodes in ad hoc and peer-to-peer networks or autonomous systems on the Internet), this poses an incentive problem; the intermediate nodes may incur significant communication and computation costs in the forwarding of packets without deriving any direct benefit from doing so. Consequently, a *rational* (i.e., utility maximizing) intermediate node may choose to

forward packets at a low priority or not forward the packets at all. This rational behavior may lead to suboptimal system performance.

The endpoints can provide incentives, e.g., in the form of payments, to encourage the intermediate nodes to forward their packets. However, the actions of the intermediate nodes are often *hidden* from the endpoints. In many cases, the endpoints can only observe whether or not the packet has reached the destination, and cannot attribute failure to a specific node on the path. Even if some form of monitoring mechanism allows them to pinpoint the location of the failure, they may still be unable to attribute the cause of failure to either the deliberate action of the intermediate node, or to some external factors beyond the control of the intermediate node, such as network congestion, channel interference, or data corruption.

The problem of hidden action is hardly unique to networks. Also known as *moral hazard*, this problem has long been of interest in the economics literature concerning information asymmetry, incentive and contract theory, and agency theory. We follow this literature by formalizing the problem as a *principal-agent* model, where multiple agents making sequential hidden actions [17, 27].

Our results are threefold. First, we show that it is possible to design contracts to induce cooperation when intermediate nodes can choose to forward or drop packets, as well as when the nodes can choose to forward packets with different levels of quality of service. If the path and transit costs are known prior to transmission, the principal achieves first best solution, and can implement the contracts either directly with each intermediate node or recursively through the network (each node making a contract with the following node) without any loss in utility. Second, we find that introducing per-hop monitoring has no impact on the principal's expected utility in equilibrium. For a principal who wishes to induce an equilibrium in which all intermediate nodes cooperate, its expected total payment is the same with or without monitoring. However, monitoring provides a *dominant strategy equilibrium*, which is a stronger solution concept than the *Nash equilibrium* achievable in the absence of monitoring. Third, we show that in the absence of a priori information about transit costs on the packet forwarding path, it is possible to generalize existing mechanisms to overcome scenarios that involve both hidden-information and hidden-action. In these scenarios, the principal pays a premium compared to scenarios with known transit costs.

2. BASELINE MODEL

We consider a principal-agent model, where the principal is a pair of communication endpoints who wish to communicate over a multi-hop network, and the agents are the intermediate nodes capable of forwarding packets between the endpoints. The principal (who in practice can be either the sender, the receiver, or

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'05, June 5–8, 2005, Vancouver, British Columbia, Canada.
Copyright 2005 ACM 1-59593-049-3/05/0006 ...\$5.00.

both) makes individual *take-it-or-leave-it* offers (*contracts*) to the agents. If the contracts are accepted, the agents choose their actions sequentially to maximize their expected payoffs based on the payment schedule of the contract. When necessary, agents can in turn make subsequent *take-it-or-leave-it* offers to their downstream agents.

We assume that all participants are risk neutral and that standard assumptions about the global observability of the final outcome and the enforceability of payments by guaranteeing parties hold.

For simplicity, we assume that each agent has only two possible actions; one involving significant effort and one involving little effort. We denote the action choice of agent i by $a_i \in \{0, 1\}$, where $a_i = 0$ and $a_i = 1$ stand for the low-effort and high-effort actions, respectively. Each action is associated with a cost (to the agent) $C(a_i)$, and we assume:

$$C(a_i = 1) > C(a_i = 0)$$

At this stage, we assume that all nodes have the same $C(a_i)$ for presentation clarity, but we relax this assumption later. Without loss of generality we normalize the $C(a_i = 0)$ to be zero, and denote the high-effort cost by c , so $C(a_i = 0) = 0$ and $C(a_i = 1) = c$.

The utility of agent i , denoted by u_i , is a function of the payment it receives from the principal (s_i), the action it takes (a_i), and the cost it incurs (c_i), as follows:

$$u_i(s_i, c_i, a_i) = s_i - a_i c_i$$

The outcome is denoted by $x \in \{x^G, x^B\}$, where x^G stands for the ‘‘Good’’ outcome in which the packet reaches the destination, and x^B stands for the ‘‘Bad’’ outcome in which the packet is dropped before it reaches the destination. The outcome is a function of the vector of actions taken by the agents on the path, $a = (a_1, \dots, a_n) \in \{0, 1\}^n$, and the loss rate on the channels, k . The benefit of the sender from the outcome is denoted by $w(x)$, where:

$$w(x^G) = w^G; \text{ and } w(x^B) = w^B = 0$$

The utility of the sender is consequently:

$$u(x, S) = w(x) - S$$

$$\text{where: } S = \sum_{i=1}^n s_i$$

A sender who wishes to induce an equilibrium in which all nodes engage in the high-effort action needs to satisfy two constraints for each agent i :

- (IR) Individual rationality (participation constraint)¹: the expected utility from participation should (weakly) exceed its reservation utility (which we normalize to 0).
- (IC) Incentive compatibility: the expected utility from exerting high-effort should (weakly) exceed its expected utility from exerting low-effort.

In some network scenarios, the topology and costs are common knowledge. That is, the sender knows in advance the path that its packet will take and the costs on that path. In other routing scenarios, the sender does not have this *a priori* information. We show that our model can be applied to both scenarios with known and unknown topologies and costs, and highlight the implications of each scenario in the context of contracts. We also distinguish between direct contracts, where the principal signs an individual contract

¹We use the notion of *ex ante* individual rationality, in which the agents choose to participate before they know the state of the system.




Figure 1: Multi-hop path from sender to destination.




Figure 2: Structure of the multihop routing game under known topology and transit costs.

with each node, and recursive contracts, where each node enters a contractual relationship with its downstream node.

The remainder of this paper is organized as follows. In Section 3 we consider agents who decide whether to drop or forward packets with and without monitoring when the transit costs are common knowledge. In Section 4, we extend the model to scenarios with unknown transit costs. In Section 5, we distinguish between recursive and direct contracts and discuss their relationship. In Section 6, we show that the model applies to scenarios in which agents choose between different levels of quality of service. We consider Internet routing as a case study in Section 7. In Section 8 we present related work, and Section 9 concludes the paper.

3. KNOWN TRANSIT COSTS

In this section we analyze scenarios in which the principal knows in advance the nodes on the path to the destination and their costs, as shown in figure 1. We consider agents who decide whether to drop or forward packets, and distinguish between scenarios with and without monitoring.

3.1 Drop versus Forward without Monitoring

In this scenario, the agents decide whether to drop ($a = 0$) or forward ($a = 1$) packets. The principal uses no monitoring to observe per-hop outcomes. Consequently, the principal makes the payment schedule to each agent contingent on the final outcome, x , as follows:

$$s_i(x) = (s_i^B, s_i^G)$$

where:

$$s_i^B = s_i(x = x^B)$$

$$s_i^G = s_i(x = x^G)$$

The timeline of this scenario is shown in figure 2. Given a per-hop loss rate of k , we can express the probability that a packet is successfully delivered from node i to its successor $i + 1$ as:

$$Pr(x_{i \rightarrow i+1}^G | a_i) = (1 - k)a_i \quad (1)$$

where $x_{i \rightarrow j}^G$ denotes a successful transmission from node i to j .

PROPOSITION 3.1. *Under the optimal contract that induces high-effort behavior from all intermediate nodes in the Nash Equi-*

librium², the expected payment to each node is the same as its expected cost, with the following payment schedule:

$$s_i^B = s_i(x = x^B) = 0 \quad (2)$$

$$s_i^G = s_i(x = x^G) = \frac{c}{(1-k)^{n-i+1}} \quad (3)$$

PROOF. The principal needs to satisfy the *IC* and *IR* constraints for each agent i , which can be expressed as follows:

$$(IC) Pr(x^G | a_{j \geq i} = 1) s_i^G + Pr(x^B | a_{j \geq i} = 1) s_i^B - c \geq \\ Pr(x^G | a_i = 0, a_{j > i} = 1) s_i^G + Pr(x^B | a_i = 0, a_{j > i} = 1) s_i^B \quad (4)$$

This constraint says that the expected utility from forwarding is greater than or equal to its expected utility from dropping, if all subsequent nodes forward as well.

$$(IR) Pr(x_{S \rightarrow i}^G | a_{j < i} = 1) (Pr(x^G | a_{j \geq i} = 1) s_i^G + \\ Pr(x^B | a_{j \geq i} = 1) s_i^B - c) + Pr(x_{S \rightarrow i}^B | a_{j < i} = 1) s_i^B \geq 0 \quad (5)$$

This constraint says that the expected utility from participating is greater than or equal to zero (reservation utility), if all other nodes forward.

The above constraints can be expressed as follows, based on Eq. 1:

$$(IC) : (1-k)^{n-i+1} s_i^G + (1 - (1-k)^{n-i+1}) s_i^B - c \geq s_i^B \\ (IR) : (1-k)^i ((1-k)^{n-i+1} s_i^G + (1 - (1-k)^{n-i+1}) s_i^B - c) + \\ (1 - (1-k)^i) s_i^B \geq 0$$

It is a standard result that both constraints bind at the optimal contract (see [23]). Solving the two equations, we obtain the solution that is presented in Eqs. 2 and 3.

We next prove that the expected payment to a node equals its expected cost in equilibrium. The expected cost of node i is its transit cost multiplied by the probability that it faces this cost (i.e., the probability that the packet reaches node i), which is: $(1-k)^i c$. The expected payment that node i receives is:

$$Pr(x^G) s_i^G + Pr(x^B) s_i^B = (1-k)^{n+1} \frac{c}{(1-k)^{n-i+1}} \\ = (1-k)^i c \quad \square$$

Note that the expected payment to a node decreases as the node gets closer to the destination due to the asymmetric distribution of risk. The closer the node is to the destination, the lower the probability that a packet will fail to reach the destination, resulting in the low payment being made to the node.

The expected payment by the principal is:

$$E[S] = (1-k)^{n+1} \sum_{i=1}^n s_i^G + (1 - (1-k)^{n+1}) \sum_{i=1}^n s_i^B \\ = (1-k)^{n+1} \sum_{i=1}^n \frac{c_i}{(1-k)^{n-i+1}} \quad (6)$$

The expected payment made by the principal depends not only on the total cost, but also the number of nodes on the path.

PROPOSITION 3.2. *Given two paths with respective lengths of n_1 and n_2 hops, per-hop transit costs of c_1 and c_2 , and per-hop loss rates of k_1 and k_2 , such that:*

²Since transit nodes perform actions sequentially, this is really a subgame-perfect equilibrium (SPE), but we will refer to it as Nash equilibrium in the remainder of the paper.




Figure 3: Two paths of equal total costs but different lengths and individual costs.

- $c_1 n_1 = c_2 n_2$ (equal total cost)
- $(1-k_1)^{n_1+1} = (1-k_2)^{n_2+1}$ (equal expected benefit)
- $n_1 < n_2$ (path 1 is shorter than path 2)

the expected total payment made by the principal is lower on the shorter path.

PROOF. The expected payment in path j is

$$E[S]_j = \sum_{i=1}^{n_j} c_j (1-k_j)^i = c_j (1-k_j) \frac{1 - (1-k_j)^{n_j}}{k_j}$$

So, we have to show that:

$$c_1 (1-k_1) \frac{1 - (1-k_1)^{n_1}}{k_1} > c_2 (1-k_2) \frac{1 - (1-k_2)^{n_2}}{k_2}$$

Let $M = c_1 n_1 = c_2 n_2$ and $N = (1-k_1)^{n_1+1} = (1-k_2)^{n_2+1}$. We have to show that

$$\frac{MN \frac{1}{n_1+1} (1 - N \frac{1}{n_1+1})}{n_1 (1 - N \frac{1}{n_1+1})} < \frac{MN \frac{1}{n_2+1} (1 - N \frac{1}{n_2+1})}{n_2 (1 - N \frac{1}{n_2+1})} \quad (7)$$

Let

$$f = \frac{N \frac{1}{n+1} (1 - N \frac{1}{n+1})}{n (1 - N \frac{1}{n+1})}$$

Then, it is enough to show that f is monotonically increasing in n

$$\frac{\partial f}{\partial n} = \frac{g(N, n)}{h(N, n)}$$

where:

$$g(N, n) = -((\ln(N)n - (n+1)^2)(N \frac{1}{n+1}) \\ - N \frac{n+2}{n+1}) - (n+1)^2 (N + N \frac{1}{n+1})$$

and

$$h(N, n) = (n+1)^2 n^2 (-1 + N \frac{1}{n+1})^2$$

but $h(N, n) > 0 \forall N, n$, therefore, it is enough to show that $g(N, n) > 0$. Because $N \in (0, 1)$: (i) $\ln(N) < 0$, and (ii) $N \frac{1}{n+1} > N \frac{n+2}{n+1}$. Therefore, $g(N, n) > 0 \forall N, n$. \square

This means that, *ceteris paribus*, shorter paths should always be preferred over longer ones.

For example, consider the two topologies presented in Figure 3. While the paths are of equal total cost, the total expected payment by the principal is different. Based on Eqs. 2 and 3, the expected total payment for the top path is:

$$E[S] = Pr(x^G) (s_A^G + s_B^G) \\ = \left(\frac{c_1}{(1-k_1)^2} + \frac{c_1}{1-k_1} \right) (1-k_1)^3 \quad (8)$$

while the expect total payment for the bottom path is:

$$E[S] = Pr(x^G)(s_A^G + s_B^G + s_C^G) \\ = \left(\frac{c_2}{(1-k_2)^3} + \frac{c_2}{(1-k_2)^2} + \frac{c_2}{1-k_2}\right)(1-k_2)^4$$

For $n_1 = 2, c_1 = 1.5, k_1 = 0.5, n_2 = 3, c_2 = 1, k_2 = 0.405$, we have equal total cost and equal expected benefit, but $E[S]_1 = 0.948$ and $E[S]_2 = 1.313$.

3.2 Drop versus Forward with Monitoring

Suppose the principal obtains per-hop monitoring information.³ Per-hop information broadens the set of mechanisms the principal can use. For example, the principal can make the payment schedule contingent on arrival to the next hop instead of arrival to the final destination. Can such information be of use to a principal wishing to induce an equilibrium in which all intermediate nodes forward the packet?

PROPOSITION 3.3. *In the drop versus forward model, the principal derives the same expected utility whether it obtains per-hop monitoring information or not.*

PROOF. The proof to this proposition is already implied in the findings of the previous section. We found that in the absence of per-hop information, the expected cost of each intermediate node equals its expected payment. In order to satisfy the *IR* constraint, it is essential to pay each intermediate node an expected amount of at least its expected cost; otherwise, the node would be better-off not participating. Therefore, no other payment scheme can reduce the expected payment from the principal to the intermediate nodes. In addition, if all nodes are incentivized to forward packets, the probability that the packet reaches the destination is the same in both scenarios, thus the expected benefit of the principal is the same. Indeed, we have found that even in the absence of per-hop monitoring information, the principal achieves first-best solution. \square

To convince the reader that this is indeed the case, we provide an example of a mechanism that conditions payments on arrival to the next hop. This is possible only if per-hop monitoring information is provided. In the new mechanism, the principal makes the payment schedule contingent on whether the packet has reached the next hop or not. That is, the payment to node i is s_i^G if the packet has reached node $i+1$, and s_i^B otherwise. We assume costless monitoring, giving us the best case scenario for the use of monitoring. As before, we consider a principal who wishes to induce an equilibrium in which all intermediate nodes forward the packet.

The expected utility of the principal is the difference between its expected benefit and its expected payment. Because the expected benefit when all nodes forward is the same under both scenarios, we only need to show that the expected total payment is identical as well. Under the monitoring mechanism, the principal has to satisfy the following constraints:

$$(IC) Pr(x_{i \rightarrow i+1}^G | a_i = 1) s^G + Pr(x_{i \rightarrow i+1}^B | a_i = 1) s^B - c \geq \\ Pr(x_{i \rightarrow i+1}^G | a_i = 0) s^G + Pr(x_{i \rightarrow i+1}^B | a_i = 0) s^B \quad (9)$$

$$(IR) Pr(x_{S \rightarrow i}^G | a_j < i = 1) (Pr(x_{i \rightarrow i+1}^G | a_i = 1) s^G \\ + Pr(x_{i \rightarrow i+1}^B | a_i = 1) s^B - c) \geq 0 \quad (10)$$

³For a recent proposal of an accountability framework that provides such monitoring information see [4].

These constraints can be expressed as follows:

$$(IC) : (1-k)s^G + ks^B - c \geq s^0 \\ (IR) : (1-k)^i((1-k)s^G + ks^B - c) \geq 0$$

The two constraints bind at the optimal contract as before, and we get the following payment schedule:

$$s^B = 0 \\ s^G = \frac{c}{1-k}$$

The expected total payment under this scenario is:

$$E[S] = \sum_{i=1}^n ((1-k)^i (s^B + (i-1)s^G)) + (1-k)^{n+1} n s^G \\ = (1-k)^{n+1} \sum_{i=1}^n \frac{c_i}{(1-k)^{n-i+1}}$$

as in the scenario without monitoring (see Equation 6.)

While the expected total payment is the same with or without monitoring, there are some differences between the two scenarios. First, the payment structure is different. If no per-hop monitoring is used, the payment to each node depends on its location (i). In contrast, monitoring provides us with n identical contracts.

Second, the *solution concept* used is different. If no monitoring is used, the strategy profile of $a_i = 1 \forall i$ is a *Nash equilibrium*, which means that no agent has an incentive to deviate *unilaterally* from the strategy profile. In contrast, with the use of monitoring, the action chosen by node i is independent of the other agents' forwarding behavior. Therefore, monitoring provides us with *dominant strategy equilibrium*, which is a stronger solution concept than Nash equilibrium. [15], [16] discuss the appropriateness of different solution concepts in the context of online environments.

4. UNKNOWN TRANSIT COSTS

In certain network settings, the transit costs of nodes along the forwarding path may not be common knowledge, i.e., there exists the problem of *hidden information*. In this section, we address the following questions:

1. Is it possible to design contracts that induce cooperative behavior in the presence of both hidden-action and hidden-information?
2. What is the principal's loss due to the lack of knowledge of the transit costs?

In hidden-information problems, the principal employs mechanisms to induce truthful revelation of private information from the agents. In the *routing game*, the principal wishes to extract transit cost information from the network routers in order to determine the lowest cost path (LCP) for a given source-destination pair. The network routers act strategically and declare transit costs to maximize their profit. Mechanisms that have been proposed in the literature for the routing game [24, 13] assume that once the transit costs have been obtained, and the LCP has been determined, the nodes on the LCP obediently forward all packets, and that there is no loss in the network, i.e., $k = 0$. In this section, we consider both hidden information and hidden action, and generalize these mechanisms to induce both truth revelation and high-effort action in equilibrium, where nodes transmit over a lossy communication channel, i.e., $k \geq 0$.

4.1 VCG Mechanism

In their seminal paper [24], Nisan and Ronen present a *VCG* mechanism that induces truthful revelation of transit costs by edges




Figure 4: Game structure for $FPSS$, where only hidden-information is considered.




Figure 5: Game structure for $FPSS'$, where both hidden-information and hidden-action are considered.

in a biconnected network, such that lowest cost paths can be chosen. Like all VCG mechanisms, it is a *strategyproof* mechanism, meaning that it induces truthful revelation in a *dominant strategy* equilibrium. In [13] ($FPSS$), Feigenbaum et al. slightly modify the model to have the routers as the selfish agents instead of the edges, and present a distributed algorithm that computes the VCG payments. The timeline of the $FPSS$ game is presented in figure 4. Under $FPSS$, transit nodes keep track of the amount of traffic routed through them via counters, and payments are periodically transferred from the principals to the transit nodes based on the counter values. $FPSS$ assumes that transit nodes are obedient in packet forwarding behavior, and will not update the counters without exerting high effort in packet forwarding.

In this section, we present $FPSS'$, which generalizes $FPSS$ to operate in an environment with lossy communication channels (i.e., $k \geq 0$) and strategic behavior in terms of packet forwarding. We will show that $FPSS'$ induces an equilibrium in which all nodes *truthfully* reveal their transit costs and forward packets if they are on the LCP. Figure 5 presents the timeline of $FPSS'$. In the first stage, the sender declares two payment functions, (s_i^G, s_i^B) , that will be paid upon success or failure of packet delivery. Given these payments, nodes have incentive to reveal their costs truthfully, and later to forward packets. Payments are transferred based on the final outcome.

In $FPSS'$, each node i submits a bid b_i , which is its reported transit cost. Node i is said to be truthful if $b_i = c_i$. We write b for the vector (b_1, \dots, b_n) of bids submitted by all transit nodes. Let $I_i(b)$ be the indicator function for the LCP given the bid vector b such that

$$I_i(b) = \begin{cases} 1 & \text{if } i \text{ is on the LCP;} \\ 0 & \text{otherwise.} \end{cases}$$

Following $FPSS$ [13], the payment received by node i at equi-

librium is:

$$\begin{aligned} p_i &= b_i I_i(b) + \left[\sum_r I_r(b)^i \infty b_r - \sum_r I_r(b) b_r \right] \\ &= \sum_r I_r(b)^i \infty b_r - \sum_{r \neq i} I_r(b) b_r \end{aligned} \quad (11)$$

where the expression $b^i x$ means that $(b^i x)_j = c_j$ for all $j \neq i$, and $(b^i x)_i = x$.

In $FPSS'$, we compute s_i^B and s_i^G as a function of p_i , k , and n . First, we recognize that s_i^B must be less than or equal to zero in order for the true LCP to be chosen. Otherwise, strategic nodes may have an incentive to report extremely small costs to mislead the principal into believing that they are on the LCP. Then, these nodes can drop any packets they receive, incur zero transit cost, collect a payment of $s_i^B > 0$, and earn positive profit.

PROPOSITION 4.1. *Let the payments of $FPSS'$ be:*

$$\begin{aligned} s_i^B &= 0 \\ s_i^G &= \frac{p_i}{(1-k)^{n-i+1}} \end{aligned}$$

Then, $FPSS'$ has a Nash equilibrium in which all nodes truthfully reveal their transit costs and all nodes on the LCP forward packets.

PROOF. In order to prove the proposition above, we have to show that nodes have no incentive to engage in the following misbehaviors:

1. truthfully reveal cost but drop packet,
2. lie about cost and forward packet,
3. lie about cost and drop packet.

If all nodes truthfully reveal their costs and forward packets, the expected utility of node i on the LCP is:

$$\begin{aligned} E[u]_i &= Pr(x_{S \rightarrow i}^G) (E[s_i] - c_i) + Pr(x_{S \rightarrow i}^B) s_i^B \\ &= (1-k)^i \left((1-k)^{n-i+1} s_i^G + (1 - (1-k)^{n-i+1}) s_i^B - c_i \right) \\ &\quad + (1 - (1-k)^i) s_i^B \\ &= (1-k)^i (1-k)^{n-i+1} \frac{p_i}{(1-k)^{n-i+1}} - (1-k)^i c_i \\ &= (1-k)^i (p_i - c_i) \\ &\geq 0 \end{aligned} \quad (12)$$

The last inequality is derived from the fact that $FPSS$ is a truthful mechanism, thus $p_i \geq c_i$. The expected utility of a node not on the LCP is 0.

A node that drops a packet receives $s_i^B = 0$, which is smaller than or equal to $E[u]_i$ for $i \in LCP$ and equals $E[u]_i$ for $i \notin LCP$. Therefore, nodes cannot gain utility from misbehaviors (1) or (3). We next show that nodes cannot gain utility from misbehavior (2).

1. if $i \in LCP$, $E[u]_i > 0$.
 - (a) if it reports $b_i > c_i$:
 - i. if $b_i < \sum_r I_r(b)^i \infty b_r - \sum_{r \neq i} I_r(b) b_r$, it is still on the LCP, and since the payment is independent of b_i , its utility does not change.
 - ii. if $b_i > \sum_r I_r(b)^i \infty b_r - \sum_{r \neq i} I_r(b) b_r$, it will not be on the LCP and obtain $E[u]_i = 0$, which is less than its expected utility if truthfully revealing its cost.

(b) if it reports $b_i < c_i$, it is still on the LCP, and since the payment is independent of b_i , its utility does not change.

2. if $i \notin LCP$, $E[u]_i = 0$.

(a) if it reports $b_i > c_i$, it remains out of the LCP, so its utility does not change.

(b) if it reports $b_i < c_i$:

i. if $b_i < \sum_r I_r(b|^\infty)b_r - \sum_{r \neq i} I_r(b)b_r$, it joins the LCP, and gains an expected utility of

$$E[u]_i = (1 - k)^i (p_i - c_i)$$

However, if $i \notin LCP$, it means that

$$c_i > \sum_r I_r(c|^\infty)c_r - \sum_{r \neq i} I_r(c)c_r$$

But if all nodes truthfully reveal their costs,

$$p_i = \sum_r I_r(c|^\infty)c_r - \sum_{r \neq i} I_r(c)c_r < c_i$$

therefore, $E[u]_i < 0$

ii. if $b_i > \sum_r I_r(b|^\infty)b_r - \sum_{r \neq i} I_r(b)b_r$, it remains out of the LCP, so its utility does not change.

Therefore, there exists an equilibrium in which all nodes truthfully reveal their transit costs and forward the received packets. \square

We note that in the hidden information only context, $FPSS$ induces truthful revelation as a dominant strategy equilibrium. In the current setting with both hidden information and hidden action, $FPSS'$ achieves a Nash equilibrium in the absence of per-hop monitoring, and a dominant strategy equilibrium in the presence of per-hop monitoring, consistent with the results in section 3 where there is hidden action only. In particular, with per-hop monitoring, the principal declares the payments s_i^B and s_i^G to each node upon failure or success of delivery to the *next* node. Given the payments $s_i^B = 0$ and $s_i^G = p_i/(1 - k)$, it is a dominant strategy for the nodes to reveal costs truthfully and forward packets.

4.2 Discussion

More generally, for any mechanism M that induces a bid vector b in equilibrium by making a payment of $p_i(b)$ to node i on the LCP, there exists a mechanism M' that induces an equilibrium with the same bid vector and packet forwarding by making a payment of:

$$s_i^B = 0$$

$$s_i^G = \frac{p_i(b)}{(1 - k)^{n-i+1}}.$$

A sketch of the proof would be as follows:

1. $I_i^M(b) = I_i^{M'}(b) \forall i$, since M' uses the same choice metric.
2. The expected utility of a LCP node is $E[u]_i = (1 - k)^i (p_i(b) - c_i) \geq 0$ if it forwards and 0 if it drops, and the expected utility of a non-LCP node is 0.
3. From 1 and 2, we get that if a node i can increase its expected utility by deviating from b_i under M' , it can also increase its utility by deviating from b_i in M , but this is in contradiction to b_i being an equilibrium in M .

4. Nodes have no incentive to drop packets since they derive an expected utility of 0 if they do.

In addition to the generalization of $FPSS$ into $FPSS'$, we can also consider the generalization of the first-price auction (FPA) mechanism, where the principal determines the LCP and pays each node on the LCP its bid, $p_i(b) = b_i$. First-price auctions achieve Nash equilibrium as opposed to dominant strategy equilibrium. Therefore, we should expect the generalization of FPA to achieve Nash equilibrium with or without monitoring.

We make two additional comments concerning this class of mechanisms. First, we find that the expected total payment made by the principal under the proposed mechanisms is

$$E[S] = \sum_{i=1}^n (1 - k)^i p_i(b)$$

and the expected benefit realized by the principal is

$$E[w] = (1 - k)^{n+1} w^G$$

where $\sum_{i=1}^n p_i$ and w^G are the expected payment and expected benefit, respectively, when only the hidden-information problem is considered. When hidden action is also taken into consideration, the generalized mechanism handles strategic forwarding behavior by conditioning payments upon the final outcome, and accounts for lossy communication channels by designing payments that reflect the distribution of risk. The difference between expected payment and benefit is *not* due to strategic forwarding behavior, but to lossy communications. Therefore, in a lossless network, we should not see any gap between expected benefits and payments, independent of strategic or non-strategic forwarding behavior.

Second, the loss to the principal due to unknown transit costs is also known as the *price of frugality*, and is an active field of research [2, 12]. This price greatly depends on the network topology and on the mechanism employed. While it is simple to characterize the principal's loss in some special cases, it is not a trivial problem in general. For example, in topologies with parallel disjoint paths from source to destination, we can prove that under first-price auctions, the loss to the principal is the difference between the cost of the shortest path and the second-shortest path, and the loss is higher under the $FPSS$ mechanism.

5. RECURSIVE CONTRACTS

In this section, we distinguish between *direct* and *recursive* contracts. In direct contracts, the principal contracts directly with each node on the path and pays it directly. In recursive payment, the principal contracts with the first node on the path, which in turn contracts with the second, and so on, such that each node contracts with its downstream node and makes the payment based on the final result, as demonstrated in figure 6.

With direct payments, the principal needs to know the identity and cost of each node on the path and to have some communication channel with the node. With recursive payments, every node needs to communicate only with its downstream node. Several questions arise in this context:

- What knowledge should the principal have in order to induce cooperative behavior through recursive contracts?
- What should be the structure of recursive contracts that induce cooperative behavior?
- What is the relation between the total expected payment under direct and recursive contracts?
- Is it possible to design recursive contracts in scenarios of unknown transit costs?




Figure 6: Structure of the multihop routing game under known topology and recursive contracts.

In order to answer the questions outlined above, we look at the *IR* and *IC* constraints that the principal needs to satisfy when contracting with the first node on the path. When the principal designs a contract with the first node, he should take into account the incentives that the first node should provide to the second node, and so on all the way to the destination.

For example, consider the topology given in figure 3 (a). When the principal comes to design a contract with node *A*, he needs to consider the subsequent contract that *A* should sign with *B*, which should satisfy the following constraints.

$$(IR) : Pr(x_{A \rightarrow B}^G | a_A = 1)(E[s_{a_B = 1}] - c) + Pr(x_{A \rightarrow B}^B | a_A = 1)s_{A \rightarrow B}^B \geq 0$$

$$(IC) : E[s_{a_B = 1}] - c \geq E[s_{a_B = 0}]$$

where:

$$E[s_{a_B = 1}] = Pr(x_{B \rightarrow D}^G | a_B = 1)s_{A \rightarrow B}^G + Pr(x_{B \rightarrow D}^B | a_B = 1)s_{A \rightarrow B}^B$$

and

$$E[s_{a_B = 0}] = Pr(x_{B \rightarrow D}^G | a_B = 0)s_{A \rightarrow B}^G + Pr(x_{B \rightarrow D}^B | a_B = 0)s_{A \rightarrow B}^B$$

These (binding) constraints yield the values of $s_{A \rightarrow B}^B$ and $s_{A \rightarrow B}^G$:

$$s_{A \rightarrow B}^B = 0 \\ s_{A \rightarrow B}^G = c / (1 - k)$$

Based on these values, *S* can express the constraints it should satisfy in a contract with *A*.

$$(IR) : Pr(x_{S \rightarrow A}^G | a_S = 1)(E[s_{S \rightarrow A} - s_{A \rightarrow B} | a_i = 1 \forall i] - c) + Pr(x_{S \rightarrow A}^B | a_S = 1)s_{S \rightarrow A}^B \geq 0$$

$$(IC) : E[s_{S \rightarrow A} - s_{A \rightarrow B} | a_i = 1 \forall i] - c \geq E[s_{S \rightarrow A} - s_{A \rightarrow B} | a_A = 0, a_B = 1]$$

where:

$$E[s_{S \rightarrow A} - s_{A \rightarrow B} | a_i = 1 \forall i] = Pr(x_{A \rightarrow D}^G | a_i = 1 \forall i)(s_{S \rightarrow A}^G - s_{A \rightarrow B}^G) + Pr(x_{A \rightarrow D}^B | a_i = 1 \forall i)(s_{S \rightarrow A}^B - s_{A \rightarrow B}^B)$$

and

$$E[s_{S \rightarrow A} - s_{A \rightarrow B} | a_A = 0, a_B = 1] = Pr(x_{A \rightarrow D}^G | a_A = 0, a_B = 1)(s_{S \rightarrow A}^G - s_{A \rightarrow B}^G) + Pr(x_{A \rightarrow D}^B | a_A = 0, a_B = 1)(s_{S \rightarrow A}^B - s_{A \rightarrow B}^B)$$

Solving for $s_{S \rightarrow A}^B$ and $s_{S \rightarrow A}^G$, we get:

$$s_{S \rightarrow A}^B = 0 \\ s_{S \rightarrow A}^G = \frac{c(2 - k)}{1 - 2k + k^2}$$

The expected total payment is

$$E[S] = s_{S \rightarrow A}^G Pr(x_{S \rightarrow D}^G) + s_{S \rightarrow A}^B Pr(x_{S \rightarrow D}^B) = c(2 - k)(1 - k) \quad (13)$$

which is equal to the expected total payment under direct contracts (see Eq. 8).

PROPOSITION 5.1. *The expected total payments by the principal under direct and recursive contracts are equal.*

PROOF. In order to calculate the expected total payment, we have to find the payment to the first node on the path that will induce appropriate behavior. Because $s_i^B = 0$ in the drop / forward model, both constraints can be reduced to:

$$Pr(x_{i \rightarrow R}^G | a_j = 1 \forall j)(s_i^G - s_{i+1}^G) - c_i = 0 \\ \Leftrightarrow (1 - k)^{n-i+1}(s_i^G - s_{i+1}^G) - c_i = 0$$

which yields, for all $1 \leq i \leq n$:

$$s_i^G = \frac{c_i}{(1 - k)^{n-i+1}} + s_{i+1}^G$$

Thus,

$$s_n^G = \frac{c_n}{1 - k}$$

$$s_{n-1}^G = \frac{c_{n-1}}{(1 - k)^2} + s_n^G = \frac{c_{n-1}}{(1 - k)^2} + \frac{c_n}{1 - k}$$

...

$$s_1^G = \frac{c_1}{(1 - k)^n} + s_2^G = \dots = \sum_{i=1}^n \frac{c_i}{(1 - k)^i}$$

and the expected total payment is

$$E[S] = (1 - k)^{n+1} s_1^G = (1 - k)^{n+1} \sum_{i=1}^n \frac{c_i}{(1 - k)^{n-i+1}}$$

which equals the total expected payment in direct payments, as expressed in Eq. 6. \square

Because the payment is contingent on the final outcome, and the expected payment to a node equals its expected cost, nodes have no incentive to offer their downstream nodes lower payment than necessary, since if they do, their downstream nodes will not forward the packet.

What information should the principal possess in order to implement recursive contracts? Like in direct payments, the expected payment is not affected solely by the total payment on the path, but also by the topology. Therefore, while the principal only needs to communicate with the first node on the forwarding path and does not have to know the identities of the other nodes, it still needs to know the number of nodes on the path and their individual transit costs.

Finally, is it possible to design recursive contracts under unknown transit costs, and, if so, what should be the structure of such contracts? Suppose the principal has implemented the distributed algorithm that calculates the necessary payments, p_i for truthful

revelation, would the following payment schedule to the first node induce cooperative behavior?

$$s_1^B = 0$$

$$s_1^G = \sum_{i=1}^n \frac{p_i}{(1-k)^i}$$

The answer is not clear. Unlike contracts in known transit costs, the expected payment to a node usually exceeds its expected cost. Therefore, transit nodes may not have the appropriate incentive to follow the principal's guarantee during the payment phase. For example, in *FPSS'*, the principal guarantees to pay each node an expected payment of $p_i > c_i$. We assume that payments are enforceable if made by the same entity that pledge to pay. However, in the case of recursive contracts, the entity that pledges to pay in the cost discovery stage (the principal) is not the same as the entity that defines and executes the payments in the forwarding stage (the transit nodes). Transit nodes, who design the contracts in the second stage, know that their downstream nodes will forward the packet as long as the expected payment exceeds the expected cost, even if it is less than the promised amount. Thus, every node has incentive to offer lower payments than promised and keep the profit. Transit nodes, who know this is a plausible scenario, may no longer truthfully reveal their cost. Therefore, while recursive contracts under known transit costs are strategically equivalent to direct contracts, it is not clear whether this is the case under unknown transit costs.

6. HIGH-QUALITY VERSUS LOW-QUALITY FORWARDING

So far, we have considered the agents' strategy space to be limited to the drop ($a = 0$) and forward ($a = 1$) actions. In this section, we consider a variation of the model where the agents choose between providing a low-quality service ($a = 0$) and a high-quality service ($a = 1$).

This may correspond to a service-differentiated service model where packets are forwarded on a best-effort or a priority basis [6]. In contrast to drop versus forward, a packet may still reach the next hop (albeit with a lower probability) even if the low-effort action is taken.

As a second example, consider the practice of hot-potato routing in inter-domain routing of today's Internet. Individual autonomous systems (AS's) can either adopt hot-potato routing or *early exit* routing ($a = 0$), where a packet is handed off to the downstream AS at the first possible exit, or *late exit* routing ($a = 1$), where an AS carries the packet longer than it needs to, handing off the packet at an exit closer to the destination. In the absence of explicit incentives, it is not surprising that AS's choose hot-potato routing to minimize their costs, even though it leads to suboptimal routes [28, 29].

In both examples, in the absence of contracts, a rational node would exert low-effort, resulting in lower performance. Nevertheless, this behavior can be avoided with an appropriate design of contracts.

Formally, the probability that a packet successfully gets from node i to node $i + 1$ is:

$$Pr(x_{i \rightarrow i+1}^G | a_i) = 1 - (k - qa_i) \quad (14)$$

where: $q \in (0, 1]$ and $k \in (q, 1]$

In the drop versus forward model, a low-effort action by any node results in a delivery failure. In contrast, a node in the high/low scenario may exert low-effort and hope to *free-ride* on the high-effort level exerted by the other agents.

PROPOSITION 6.1. *In the high-quality versus low-quality forwarding model, where transit costs are common knowledge, the principal derives the same expected utility whether it obtains per-hop monitoring information or not.*

PROOF. The *IC* and *IR* constraints are the same as specified in the proof of proposition 3.1, but their values change, based on Eq. 14 to reflect the different model:

$$(IC) : (1-k+q)^{n-i+1} s_i^G + (1-(1-k+q)^{n-i+1}) s_i^B - c \geq (1-k)(1-k+q)^{n-i} s_i^G + (1-(1-k)(1-k+q)^{n-i}) s_i^B$$

$$(IR) : (1-k+q)^i ((1-k+q)^{n-i+1} s_i^G + (1-(1-k+q)^{n-i+1}) s_i^B - c) + (1-(1-k+q)^i) s_i^B \geq 0$$

For this set of constraints, we obtain the following solution:

$$s_i^B = \frac{(1-k+q)^i c(k-1)}{q} \quad (15)$$

$$s_i^G = \frac{(1-k+q)^i c(k-1 + (1-k+q)^{-n})}{q} \quad (16)$$

We observe that in this version, both the high and the low payments depend on i . If monitoring is used, we obtain the following constraints:

$$(IC) : (1-k+q) s_i^G + (k-q) s_i^B - c \geq (1-k) s_i^G + (k) s_i^B$$

$$(IR) : (1-k+q)^i ((1-k+q) s_i^G + (k-q) s_i^B - c) \geq 0$$

and we get the solution:

$$s_i^B = \frac{c(k-1)}{q}$$

$$s_i^G = \frac{ck}{q}$$

The expected payment by the principal with or without forwarding is the same, and equals:

$$E[S] = \frac{c(1-k+q)(1-(1-k+q)^n)}{k-q} \quad (17)$$

and this concludes the proof. \square

The payment structure in the high-quality versus low-quality forwarding model is different from that in the drop versus forward model. In particular, at the optimal contract, the low-outcome payment s_i^B is now less than zero. A negative payment means that the agent must pay the principal in the event that the packet fails to reach the destination. In some settings, it may be necessary to impose a *limited liability* constraint, i.e., $s_i \geq 0$. This prevents the first-best solution from being achieved.

PROPOSITION 6.2. *In the high-quality versus low-quality forwarding model, if negative payments are disallowed, the expected payment to each node exceeds its expected cost under the optimal contract.*

PROOF. The proof is a direct outcome of the following statements, which are proved above:

1. The optimal contract is the contract specified in equations 15 and 16
2. Under the optimal contract, $E[s_i]$ equals node i 's expected cost
3. Under the optimal contract, $s_i^B = \frac{(1-k+q)^i c(k-1)}{q} < 0$

Therefore, under any other contract the sender will have to compensate each node with an expected payment that is higher than its expected transit cost. \square

There is an additional difference between the two models. In drop versus forward, a principal either signs a contract with all n nodes along the path or with none. This is because a single node dropping the packet determines a failure. In contrast, in high versus low-quality forwarding, a success may occur under the low effort actions as well, and payments are used to increase the probability of success. Therefore, it may be possible for the principal to maximize its utility by contracting with only m of the n nodes along the path. While the expected outcome depends on m , it is independent of which specific m nodes are induced. At the same time, the individual expected payments decrease in i (see Eq. 16). Therefore, a principal who wishes to sign a contract with only m out of the n nodes should do so with the nodes that are closest to the destination; namely, nodes $(n - m + 1, \dots, n - 1, n)$.

Solving for the high-quality versus low-quality forwarding model with unknown transit costs is left for future work.

7. CASE STUDY: INTERNET ROUTING

We can map different deployed and proposed Internet routing schemes to the various models we have considered in this work.

Border Gateway Protocol (BGP), the current inter-domain routing protocol in the Internet, computes routes based on path vectors. Since the protocol reveals only the autonomous systems (AS's) along a route but not the cost associated to them, the current BGP routing is best characterized by lack of a priori information about transit costs. In this case, the principal (e.g., a multi-homed site or a tier-1 AS) can implement one of the mechanisms proposed in Section 4 by contracting with individual nodes on the path. Such contracts involve paying some premium over the real cost, and it is not clear whether recursive contracts can be implemented in this scenario. In addition, the current protocol does not have the infrastructure to support implementation of direct contracts between endpoints and the network.

Recently, several new architectures have been proposed in the context of the Internet to provide the principal not only with a set of paths from which it can choose (like BGP does) but also with the performance along those paths and the network topology. One approach to obtain such information is through end-to-end probing [1]. Another approach is to have the edge networks perform measurements and discover the network topology [32]. Yet another approach is to delegate the task of obtaining topology and performance information to a third-party, like in the routing-as-a-service proposal [21]. These proposals are quite different in nature, but they are common in their attempt to provide more visibility and transparency into the network. If information about topology and transit costs is obtained, the scenario is mapped to the "known transit costs" model (Section 3). In this case, first-best contracts can be achieved through individual contracts with nodes along the path. However, as we have shown in Section 5, as long as each agent can choose the next hop, the principal can gain full benefit by contracting with only the first hop (through the implementation of recursive contracts).

However, the various proposals for acquiring network topology and performance information do not deal with strategic behavior by the intermediate nodes. With the realization that the information collected may be used by the principal in subsequent contractual relationships, the intermediate nodes may behave strategically, misrepresenting their true costs to the entities that collect and aggregate such information. One recent approach that can alleviate this problem is to provide packet obituaries by having each packet to confirm its delivery or report its last successful AS hop [4]. Another approach is to have third parties like Keynote independently monitor the network performance.

8. RELATED WORK

The study of non-cooperative behavior in communication networks, and the design of incentives, has received significant attention in the context of wireless ad-hoc routing. [22] considers the problem of *malicious* behavior, where nodes respond positively to route requests but then fail to forward the actual packets. It proposes to mitigate it by detection and report mechanisms that will essentially help to route around the malicious nodes. However, rather than penalizing nodes that do not forward traffic, it bypasses the misbehaving nodes, thereby relieving their burden. Therefore, such a mechanism is not effective against *selfish* behavior.

In order to mitigate selfish behavior, some approaches [7, 8, 9] require reputation exchange between nodes, or simply first-hand observations [5]. Other approaches propose payment schemes [10, 20, 31] to encourage cooperation. [31] is the closest to our work in that it designs payment schemes in which the sender pays the intermediate nodes in order to prevent several types of selfish behavior. In their approach, nodes are supposed to send receipts to a third-party entity. We show that this type of per-hop monitoring may not be needed.

In the context of Internet routing, [4] proposes an accountability framework that provide end hosts and service providers after-the-fact audits on the fate of their packets. This proposal is part of a broader approach to provide end hosts with greater control over the path of their packets [3, 30]. If senders have transit cost information and can fully control the path of their packets, they can design contracts that yield them with first-best utility. The accountability framework proposed in [4] can serve two main goals: informing nodes of network conditions to help them make informed decision, and helping entities to establish whether individual ASs have performed their duties adequately. While such a framework can be used for the first task, we propose a different approach to the second problem without the need of per-hop auditing information.

Research in distributed algorithmic mechanism design (DAMD) has been applied to BGP routing [13, 14]. These works propose mechanisms to tackle the hidden-information problem, but ignore the problem of forwarding enforcement. Inducing desired behavior is also the objective in [26], which attempts to respond to the challenge of distributed AMD raised in [15]: if the same agents that seek to manipulate the system also run the mechanism, what prevents them from deviating from the mechanism's proposed rules to maximize their own welfare? They start with the proposed mechanism presented in [13] and use mostly auditing mechanisms to prevent deviation from the algorithm.

The focus of this work is the design of a payment scheme that provides the appropriate incentives within the context of multi-hop routing. Like other works in this field, we assume that all the accounting services are done using out-of-band mechanisms. Security issues within this context, such as node authentication or message encryption, are orthogonal to the problem presented in this paper, and can be found, for example, in [18, 19, 25].

The problem of information asymmetry and hidden-action (also known as moral hazard) is well studied in the economics literature [11, 17, 23, 27]. [17] identifies the problem of moral hazard in production teams, and shows that it is impossible to design a sharing rule which is efficient and budget-balanced. [27] shows that this task is made possible when production takes place sequentially.

9. CONCLUSIONS AND FUTURE DIRECTIONS

In this paper we show that in a multi-hop routing setting, where the actions of the intermediate nodes are hidden from the source

and/or destination, it is possible to design payment schemes to induce cooperative behavior from the intermediate nodes. We conclude that monitoring per-hop outcomes may not improve the utility of the participants or the network performance. In addition, in scenarios of unknown transit costs, it is also possible to design mechanisms that induce cooperative behavior in equilibrium, but the sender pays a premium for extracting information from the transit nodes.

Our model and results suggest several natural and intriguing research avenues:

- Consider manipulative or collusive behaviors which may arise under the proposed payment schemes.
- Analyze the feasibility of recursive contracts under hidden-information of transit costs.
- While the proposed payment schemes sustain cooperation in equilibrium, it is not a *unique* equilibrium. We plan to study under what mechanisms this strategy profile may emerge as a unique equilibrium (e.g., penalty by successor nodes).
- Consider the effect of congestion and capacity constraints on the proposed mechanisms. Our preliminary results show that when several senders compete for a single transit node's capacity, the sender with the highest demand pays a premium even if transit costs are common knowledge. The premium can be expressed as a function of the second-highest demand. In addition, if congestion affects the probability of successful delivery, a sender with a lower cost alternate path may end up with a lower utility level than his rival with a higher cost alternate path.
- Fully characterize the full-information Nash equilibrium in first price auctions, and use this characterization to derive its *overcharging* compared to truthful mechanisms.

10. ACKNOWLEDGEMENTS

We thank Hal Varian for his useful comments. This work is supported in part by the National Science Foundation under ITR awards ANI-0085879 and ANI-0331659, and Career award ANI-0133811.

11. REFERENCES

- [1] ANDERSEN, D. G., BALAKRISHNAN, H., KAASHOEK, M. F., AND MORRIS, R. Resilient Overlay Networks. In *18th ACM SOSP* (2001).
- [2] ARCHER, A., AND TARDOS, E. Frugal path mechanisms.
- [3] ARGYRAKI, K., AND CHERITON, D. Loose Source Routing as a Mechanism for Traffic Policies. In *Proceedings of SIGCOMM FDNA* (August 2004).
- [4] ARGYRAKI, K., MANIATIS, P., CHERITON, D., AND SHENKER, S. Providing Packet Obituaries. In *Third Workshop on Hot Topics in Networks (HotNets)* (November 2004).
- [5] BANSAL, S., AND BAKER, M. Observation-based cooperation enforcement in ad hoc networks. *Technical report, Stanford university* (2003).
- [6] BLAKE, S., BLACK, D., CARLSON, M., DAVIES, E., WANG, Z., AND WEISS, W. An Architecture for Differentiated Service. RFC 2475, 1998.
- [7] BUCHEGGER, S., AND BOUDEC, J.-Y. L. Performance Analysis of the CONFIDANT Protocol: Cooperation of Nodes - Fairness in Dynamic Ad Hoc Networks. In *IEEE/ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC)* (2002).
- [8] BUCHEGGER, S., AND BOUDEC, J.-Y. L. Coping with False Accusations in Misbehavior Reputation Systems For Mobile Ad Hoc Networks. In *EPFL, Technical report* (2003).
- [9] BUCHEGGER, S., AND BOUDEC, J.-Y. L. The effect of rumor spreading in reputation systems for mobile ad-hoc networks. In *WiOpt'03: Modeling and Optimization in Mobile Ad Hoc and Wireless Networks* (2003).
- [10] BUTTYAN, L., AND HUBAUX, J. Stimulating Cooperation in Self-Organizing Mobile Ad-Hoc Networks. *ACM/Kluwer Journal on Mobile Networks and Applications (MONET)* (2003).
- [11] CAILLAUD, B., AND HERMALIN, B. *Hidden Action and Incentives*. Teaching Notes. U.C. Berkeley.
- [12] ELKIND, E., SAHAI, A., AND STEIGLITZ, K. Frugality in path auctions, 2004.
- [13] FEIGENBAUM, J., PAPADIMITRIOU, C., SAMI, R., AND SHENKER, S. A BGP-based Mechanism for Lowest-Cost Routing. In *Proceedings of the ACM Symposium on Principles of Distributed Computing* (2002).
- [14] FEIGENBAUM, J., SAMI, R., AND SHENKER, S. Mechanism Design for Policy Routing. In *Yale University, Technical Report* (2003).
- [15] FEIGENBAUM, J., AND SHENKER, S. Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. In *Proceedings of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications* (2002).
- [16] FRIEDMAN, E., AND SHENKER, S. Learning and implementation on the internet. In *Manuscript. New Brunswick: Rutgers University, Department of Economics* (1997).
- [17] HOLMSTROM, B. Moral Hazard in Teams. *Bell Journal of Economics* 13 (1982), 324–340.
- [18] HU, Y., PERRIG, A., AND JOHNSON, D. Ariadne: A Secure On-Demand Routing Protocol for Ad Hoc Networks. In *Eighth Annual International Conference on Mobile Computing and Networking (Mobicom)* (2002), pp. 12–23.
- [19] HU, Y., PERRIG, A., AND JOHNSON, D. SEAD: Secure Efficient Distance Vector Routing for Mobile Ad Hoc Networks. In *4th IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)* (2002).
- [20] JAKOBSSON, M., HUBAUX, J.-P., AND BUTTYAN, L. A Micro-Payment Scheme Encouraging Collaboration in Multi-Hop Cellular Networks. In *Financial Cryptography* (2003).
- [21] LAKSHMINARAYANAN, K., STOICA, I., AND SHENKER, S. Routing as a service. In *UCB Technical Report No. UCB/CSD-04-1327* (January 2004).
- [22] MARTI, S., GIULI, T. J., LAI, K., AND BAKER, M. Mitigating Routing Misbehavior in Mobile Ad Hoc Networks. In *Proceedings of MobiCom* (2000), pp. 255–265.
- [23] MASS-COLELL, A., WHINSTON, M., AND GREEN, J. *Microeconomic Theory*. Oxford University Press, 1995.
- [24] NISAN, N., AND RONEN, A. Algorithmic Mechanism Design. In *Proceedings of the 31st Symposium on Theory of Computing* (1999).
- [25] SANZGIRI, K., DAHILL, B., LEVINE, B., SHIELDS, C., AND BELDING-ROYER, E. A Secure Routing Protocol for Ad Hoc Networks. In *International Conference on Network Protocols (ICNP)* (2002).
- [26] SHNEIDMAN, J., AND PARKES, D. C. Overcoming rational manipulation in mechanism implementation, 2004.
- [27] STRAUSS, R. Moral Hazard in Sequential Teams. *Departmental Working Paper. Free University of Berlin* (1996).
- [28] TEIXEIRA, R., GRIFFIN, T., SHAIKH, A., AND VOELKER, G. Network sensitivity to hot-potato disruptions. In *Proceedings of ACM SIGCOMM* (September 2004).
- [29] TEIXEIRA, R., SHAIKH, A., GRIFFIN, T., AND REXFORD, J. Dynamics of hot-potato routing in IP networks. In *Proceedings of ACM SIGMETRICS* (June 2004).
- [30] YANG, X. NIRA: A New Internet Routing Architecture. In *Proceedings of SIGCOMM FDNA* (August 2003).
- [31] ZHONG, S., CHEN, J., AND YANG, Y. R. Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-Hoc Networks. In *22nd Annual Joint Conference of the IEEE Computer and Communications Societies* (2003).
- [32] ZHU, D., GRITTER, M., AND CHERITON, D. Feedback-based Routing. In *Proc Hotnets-I* (2002).