# Hidden-Action in Network Routing

Michal Feldman, John Chuang, Ion Stoica, and Scott Shenker

***Abstract*— In communication networks, such as the Internet or mobile ad-hoc networks, the actions taken by intermediate nodes or links 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., selfish) intermediaries may choose to forward messages at a low priority or simply not forward messages 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 directly) and the recursive (each router contracts with the next downstream router) cases. We further show that, depending on the network topology, per-hop or per-path monitoring may not necessarily improve the utility of the principal or the social welfare of the system.**

***Index Terms*— Hidden action, moral hazard, routing, economics, contract, principal-agent, mechanism design, game theory.**

## I. Introduction

COMMUNICATION networks such as the Internet or mobile ad-hoc networks are deployed, operated and used by multiple parties with diverse economic interests. The performance of these systems is governed by interactions among self-interested parties with varying degrees of cooperation and competition. It is therefore important for the analysis and design of network systems and protocols to take economic self-interest into account. Researchers have embraced game theory as a powerful tool for understanding the effect of cooperation and for building networks in which selfish agents engage in cooperative behavior.

In this paper we study the problem of selfish behavior in network routing, or more precisely, message forwarding. Endpoints wishing to communicate over a network rely on intermediate nodes to forward messages from the sender to the receiver. In settings where the intermediaries are independent agents, e.g., individual devices in mobile ad-hoc networks, peers in peer-to-peer networks, autonomous systems in the Internet, an incentive problem arises. The intermediaries may incur significant communication and computation costs in message forwarding without deriving any direct benefit from doing so. Consequently, a *rational*, i.e., utility maximizing, intermediary may engage in one of two types of misbehaviors. First, it may misreport some private information it might possess, e.g., channel loss rate, transit cost, to maximize its profit from forwarding messages, or to avoid forwarding messages altogether. Second, it may forward the messages at a low priority or not forward the messages at all. This individually rational behavior may lead to suboptimal system performance.

Algorithmic mechanism design (AMD) [19], [7], [8] provides a mathematical framework to study the problem of *hidden information*, focusing on the design of mechanisms that extract private information from rational agents. On the other hand, the problem of *hidden action* has received limited formal treatment in the networking literature.

Various payment-based and reputation-based incentive mechanisms have been proposed to mitigate selfish behavior in message forwarding [6], [4], [20], [14], [30], [28]. However, the actions of the intermediaries are often *hidden* from the endpoints. In many cases, the endpoints can only observe whether or not a message has reached the destination, and cannot attribute failure to a specific node or link on the path. Even if some form of monitoring mechanism allows the endpoints 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 intermediary, or to some external factors beyond its control, such as network congestion, channel interference, or data corruption.

This problem of hidden action arises in a variety of network contexts. In Internet routing, neither *ex ante* nor *post hoc* information is provided about the fate of packets. While lack of assurances enables flexibility, the lack of accountability also enables IP routers to delay or drop packets for a variety of reasons without notification of the endpoints. Hidden-action also arises in mobile ad hoc networks, where devices may strategically drop packets to conserve their constrained energy resources, or in a variety of application-layer overlay networks such as distributed hash tables (DHTs) [22], [23], peer-to-peer networks for file-sharing, resilient routing [1], and anonymous routing [12]. The highly dynamic memberships and topologies of these networks make monitoring extremely difficult.

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 [17]. Holmstrom [13] has first identified the problem of moral hazard in production teams, and many subsequent works on team incentives followed [24], [21], [16], [29]. We follow this literature by formalizing the problem as a *principal-agent* model, where a principal motivates a team of strategic agents to exert costly effort on his behalf, but their actions are hidden from him.

Manuscript received July 1, 2006; revised February 15, 2007. A preliminary version of this paper was presented at the ACM conference on Electronic Commerce (ACM EC 2005).

Michal Feldman is with School of Engineering and Computer Science, The Hebrew University of Jerusalem, Israel (e-mail: mfeldman@cs.huji.ac.il).

John Chuang is with School of Information, University of California at Berkeley (e-mail: chuang@sims.berkeley.edu).

Ion Stoica and Scott Shenker are with Computer Science Division, University of California at Berkeley (e-mail: istoica@cs.berkeley.edu, shenker@icsi.berkeley.edu).

Digital Object Identifier 10.1109/JSAC.2007.070810.

By modeling the communicating endpoints as the principal and the network links as the agents, we show that it is possible to design contracts that induce an equilibrium with cooperative behavior by the rational agents. Our results are fourfold.

First, we quantify the effect of network topology on the utility loss by the principal due to hidden action. In sequential-links topologies, the principal incurs no loss due to his inability to observe individual actions. Thus, employing any kind of monitoring mechanism within the network has no impact on the principal's utility. In contrast, in parallel-links topologies, the principal incurs some loss due to hidden action, and can therefore improve his utility from obtaining per-link monitoring information. Yet, in parallel-paths topologies, per-path information already provides the principal with all the information needed to achieve the highest possible utility. In this case, per-link information yields no additional utility to the principal. However, per-link monitoring provides a *dominant strategy equilibrium*, which is a stronger solution concept than the *Nash equilibrium* achievable in the absence of monitoring.

Second, we show that it is possible to implement the contracts either directly (between the principal and each of the agents individually) or recursively (where the principal contracts with its downlink agent, who in turn contracts with its downlink agent, etc.) The relative performance of recursive contracts versus direct contracts is influenced not just by the network topology, but also by the loss rate.

Third, we show that in scenarios that are characterized by both hidden action and hidden information, appropriate design of contracts can induce a Nash equilibrium in which all agents truthfully reveal their private information and also exert the desired level of effort.

Finally, we consider a quality of service (QoS) model, where the agents can choose to forward messages at different effort levels. In this case, successful message delivery can be accomplished even if the agents do not engage in the highest effort level, but the probability of success is diminished. Then, we show that the principal incurs a utility loss due to hidden action, even in sequential-links topologies, and this can be attributed to the problem of free-riding.

## II. MODEL

We develop a principal-agent model, where the principal is a pair of communication endpoints that wish to communicate over a multi-hop network, and the agents are the links that lie on the path between the endpoints, and are capable of forwarding messages between them. We choose to model the links as agents for consistency with the literature and for a more elegant exposition, but our results can be generalized for a model where the nodes are the agents.

The principal (who in practice can be either the sender, the receiver, or both) makes individual *take-it-or-leave-it* offers (*contracts*) to the agents. If the contracts are accepted, the agents choose their actions to maximize their expected payoffs based on the payment schedule of the contract.

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:

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

We also normalize $C(a_i = 0)$ to be zero, and denote the high-effort cost by $c_i$, so $C(a_i = 0) = 0$ and $C(a_i = 1) = c$. That is, $C(a_i) = c_i a_i$.

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 graph, $a = (a_1, ..., a_n) \in \{0,1\}^n$, and the loss rate on the channels, denoted by $k$. The benefit of the sender from the outcome is denoted by $w(x)$, where:

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

The utility of the sender is consequently:

$$u(x, S) = w(x) - S$$
$$\text{where: } S = \sum_{i=1}^{n} s_i$$

The principal motivates the agents by offering a set of contracts, putting the agents on an equilibrium point of the induced game. Unless otherwise stated, we restrict attention to the case in which the principal induces all agent to exert effort [1]. In order to induce an agent $i$ to exert effort, the principal needs to satisfy the following two constraints for that agent:

**IR** Individual rationality (participation constraint)[2]: the expected utility of the agent from participation should (weakly) exceed 0.

**IC** Incentive compatibility: the expected utility of the agent from exerting high-effort should (weakly) exceed its expected utility from exerting low-effort.

In some scenarios, it is also reasonable to impose the *limited liability* (LL) constraint, which states that the agents never pay the principal. We later discuss the implications of this constraint.

Throughout the paper we deal with various network topologies, such as sequential links, parallel links and parallel paths. When two links are parallel, their actions are assumed to be taken simultaneously, and we use the Nash equilibrium solution concept. When two links are in sequence, their actions

[1] A follow-up paper [3] extends the model and studies the optimal set of agents to be contracted, given the principal's benefit ($w^G$).

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

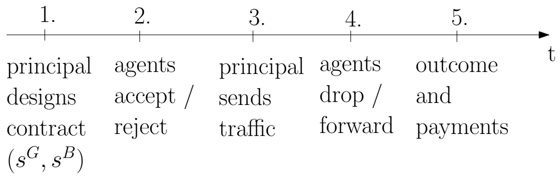


Fig. 1. Structure of the multihop routing game under known topology and transit costs.

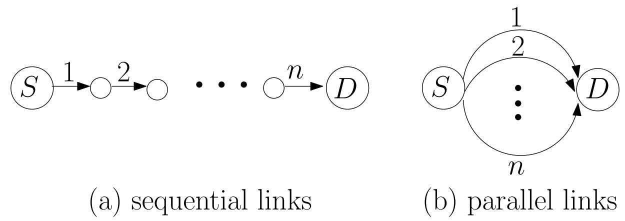


Fig. 2. Network topologies from sender to destination.

are assumed to be taken sequentially, and the corresponding solution concept is the subgame perfect equilibrium. In the remainder of this paper we refer to both equilibria simply as equilibrium points, where the specific solution concept can be inferred from the context.

In some network scenarios, the topology and costs are common knowledge. That is, the sender knows in advance the structure of the graph that connects the source and the destination, and the cost of each edge of the graph. In other routing scenarios, the sender does not have this *a priori* information. We discuss the implications of unknown topologies and costs in the context of hidden actions. We also distinguish between direct contracts, where the principal signs an individual contract with each agent, and recursive contracts, where each agent enters a contractual relationship with its downstream agent.

The proofs that do not appear in the body of the paper appear in the Appendix.

## III. Optimal Contracts Under Hidden Actions

In this section we analyze scenarios in which the topology and the costs are common knowledge. We consider agents who decide whether to drop ($a = 0$) or forward ($a = 1$) packets, and distinguish between results under different topologies.

Since the principal observes only the final outcome (i.e., whether the packet has reached the destination or not), he can make 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 1. That is, the principal designs and announces the payments of success and failure, and the agents accept or reject the contract according to the announced contract. The principal then sends the packet, the agents decide whether to drop or forward the packet, and the principal pays the agents, based on the contract, after he observes whether the packet reached the destination or not.

Given a per-hop loss rate of $k < 1$, the probability that a packet is successfully delivered along a link $i$ is given by:

$$Pr(x_i^G | a_i) = (1 - k)a_i \quad (1)$$

where $x_i^G$ denotes a successful transmission on link $i$. That is, if $a_i = 1$, the packet is successfully delivered along link $i$ with probability $1 - k$, and if $a_i = 0$, link $i$ necessarily fails. This model is used throughout the whole paper, except for Section VI, which considers the quality-of-service model. It is also helpful to denote by $x_{S\to i}^G$ a successful transmission from the source to link $i$, and by $x_{i\to D}^G$ a successful transmission from link $i$ to the destination (and similarly a failure by $x_{S\to i}^B$ and $x_{i\to D}^B$).

### A. Sequential Links

We begin with the case in which the agents are placed in a sequential path from the source node to the destination node, as shown in Figure 2(a). For presentation clarity, we first assume that all agents have the same cost function, though we relax this assumption later.

When all the agents exert effort, the probability of a successful delivery is given by:

$$Pr(x^G | \forall i \ a_i = 1) = (1 - k)^n$$

The following proposition shows that in the optimal contract that induces all agents to exert effort, the expected payment to each link equals to its expected cost. Thus, a sender who wishes to send traffic through a path of sequential links does not lose utility (in expectation) from his inability to monitor the individual actions.

*Proposition 3.1:* In sequential-links scenarios, under the optimal contract that induces a Nash equilibrium in which all the agents exert high-effort, the expected payment to each link 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 $IC$ and $IR$ constraints for each agent $i$ can be expressed as follows:

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

This constraint says that the expected utility from forwarding is not lower than its expected utility from dropping, given that all subsequent links forward as well.

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

This constraint says that the expected utility from participating is not lower than zero (his reservation utility), given that all other links forward.

Based on Eq. 1, the above constraints can be expressed as follows:

$$\textbf{IC} \quad (1-k)^{n-i+1}s_i^G + (1-(1-k)^{n-i+1})s_i^B - c \geq s_i^B \tag{6}$$

$$\textbf{IR} \quad (1-k)^i((1-k)^{n-i+1}s_i^G + (1-(1-k)^{n-i+1})s_i^B - c) + (1-(1-k)^i)s_i^B \geq 0 \tag{7}$$

In the case of sequential links, both constraints bind at the optimal contract. Solving the two equations, we obtain the solution that is presented in Eqs. 2 and 3, as required.

We next prove that the expected payment to a link equals its expected cost in equilibrium. The expected cost of link $i$ is its transit cost multiplied by the probability that it faces this cost (i.e., the probability that the message reaches link $i$), which is given by $E[c_i] = (1-k)^{i-1}c$.

The expected payment that link $i$ receives is:

$$\begin{aligned} E[s_i] &= Pr(x^G)s_i^G + Pr(x^B)s_i^B \\ &= (1-k)^n \frac{c}{(1-k)^{n-i+1}} + \\ & \quad (1-(1-k)^n) \cdot 0 \\ &= (1-k)^{i-1}c \end{aligned}$$

■

Note that the expected payment to a link decreases as the link is located closer to the destination due to the asymmetric distribution of risk. The closer the link is to the destination, the lower the probability that a packet will fail to reach the destination, resulting in the lower payment made to the link. The expected payment by the principal is:

$$\begin{aligned} E[s] &= (1-k)^n \sum_{i=1}^{n} s_i^G + (1-(1-k)^n) \sum_{i=1}^{n} s_i^B \\ &= (1-k)^n \sum_{i=1}^{n} \frac{c_i}{(1-k)^{n-i+1}} \end{aligned} \tag{8}$$

The expected payment made by the principal depends not only on the total cost on the path, but also on the length of the path (i.e., the number of agents on the path). We next study the effect of the path length on the principal's utility, assuming all other things equal. That is, we compare the principal's utility in two scenarios that have different length (number of links), but equal probability of success and equal total cost. Interestingly enough, we find that longer paths should be preferred over shorter ones.

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

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

the expected principal's utility in the longer path is higher than his expected utility in the shorter path.

### *B. Parallel Links*

In some scenarios the sender may elect to send multiple copies of the packets along disjoint paths. For example, multipath routing [11] refers to routing techniques in which multiple copies of the message are sent through multiple paths in order to increase the likelihood of reliable data delivery. Multipath routing for fault tolerance has been studied, for example, in the contexts of high-speed networks [18] and wireless networks. We begin our discussion with the simplest topology of disjoint paths, which already gives us some insights regarding the effect of the topology in our model. This topology is composed of $n$ parallel links from the source to the destination, as demonstrated in Figure 2(b). In contrast to the sequential-links topology, where agents have different information when they decide which actions to take, here, all the agents act simultaneously, based on no prior information.

We assume that the principal sends a copy of the message on each of the links, but observes only whether the delivery was successful. A successful delivery means that at least one copy has reached the destination. Upon a successful delivery, all the agents receive the high payments. Otherwise, they all get the low payments [3]. The probability of a successful delivery when all $n$ agents forward the packet is given by:

$$Pr(x^G|\forall i \; a_i = 1) = 1 - k^n$$

It turns out that unlike the sequential-links topology, the limited liability (LL) constraint has a qualitative effect on the results. Recall that the limited liability constraint states that agents never pay the principal. i.e., for any $i$, $s_i^G, s_i^B \geq 0$. In contrast, if the limited liability constraint is not imposed, the payments can take any value; in particular, $S_i^G$ may be negative.

*Proposition 3.3:* In the parallel-links scenario, under the optimal contract that induces high-effort behavior from all intermediate edges in a Nash Equilibrium, the expected payment to each link exceeds its expected cost if the limited-liability constraint is imposed, and equals to its expected cost otherwise.

This result suggests that in the parallel-links scenario, if the LL constraint is imposed, the principal pays each agent more than its cost. However, if the principal is allowed to impose penalty upon a failure, then paying each agent exactly its expected cost is possible. The intuition for this result is that when the message can reach the destination through other paths, it brings the possibility for agents to free ride on each others' efforts without bearing the risk to receive a lower payment. In this case, additional payment is required to avoid the free-riding problem. If, however, the LL constraint is not imposed and $s_i^B < 0$, it is still possible to deter agent $i$ from dropping the message with a payment schedule that satisfies $E[c_i] = E[s_i]$. In the remainder of this paper, we assume that the LL constraint is imposed, thus $s_i^B = 0$. This is a standard assumption in much of the principal-agent literature, since otherwise, the principal does not lose utility from his inability to observe individual actions. In addition, when the

[3]In this scenario, the principal might wish to induce only a subset of the agents to exert effort. For further discussion about contracting with some subset of the agents, see [3]

LL constraint is imposed, the IC constraint is always binding. That is, a payment schedule that satisfies the IC constraint also satisfies the IR constraint. Thus, in the remainder of the paper, we restrict attention to the IC constraint.

## IV. The Effect of Monitoring

In Section III we assumed that the principal can only observe the final outcome. i.e., whether the packet has reached the destination or not. This assumption might fit some communication protocols, while others may provide the principal with more information about the intermediate outcomes within the network. In this section, we study the effect of various levels of monitoring under different topologies.

One can think of several levels of monitoring. Obviously, if *individual actions* are observable, then paying each agent its cost if it exerts effort ensures cooperative behavior. If only the final outcome is observable, then the principal may need to pay each agent more than its expected cost, depending on the topology, as concluded in the previous section. Between observing individual actions and observing only the final outcome, there are various intermediate levels of monitoring. For example, in a network that is composed of multiple parallel paths (see Figure 3), there may be several possible levels of monitoring. The sender may learn only whether the packet has reached the destination or not. He might also gain information regarding the number of copies reaching the destination. Yet more specific, he may learn through which paths copies of the packet arrived. In this section we study the effect of the monitoring level on the optimal contract in different topologies.

### A. Sequential Links

Per-hop information broadens the set of mechanisms the principal can use. A natural mechanism for the principal is to make the payment schedule contingent on the individual outcome (i.e., arrival to the next hop). Can such flexibility be of use to a principal wishing to induce an equilibrium in which all intermediate edges forward the packet?

*Proposition 4.1:* In the sequential-links topology, the principal derives the same expected utility whether it obtains per-hop monitoring information or not.

This result is specific to the model in which an agent that does not exert effort fails with certainty (for a different model, see Section VI). Under this assumption, a message can reach the destination only if all the agents exert effort in forwarding, therefore, per-hop monitoring cannot help. The following example demonstrates that per-hop monitoring does not improve the principal's expected utility.

*Claim 4.2:* In sequential-links topologies, a mechanism that makes payments contingent on arrival to the next link (i.e., pays $s_i^G$ if the message reached link $i+1$, and pays $s_i^B$ otherwise) yields the same expected utility to the principal as one that makes the payments contingent on arrival to the destination.

While the expected total payment is the same with or without monitoring, the *solution concept* 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, given that the other agents exert effort. 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 a *dominant strategy* equilibrium, which is a stronger solution concept than Nash equilibrium[4].

### B. Parallel Links

In Section III we learned that the expected payment to each agent exceeds its expected cost. How would per-hop monitoring affect this result? The following proposition shows that when per-hop monitoring is available, the principal obtains a higher utility since he can achieve a contract in which each agent receives an expected payment that is equal to its expected cost. This result holds independent of the topology.

*Claim 4.3:* In any network topology, when per-hop monitoring is available, the principal achieves a contract in which he pays each agent his expected cost. The payment schedule is $s_i^B = 0$ and $s_i^G = \frac{c}{1-k}$.

### C. Parallel Paths

We next analyze the effect of monitoring in topologies with parallel paths between the source and the destination. If the sender can observe only whether the packet has reached the destination or not, the hidden-action problem entails some loss to the principal, as in the parallel-links scenario (see Section IV-B above). Obviously, monitoring the individual (per-hop) outcomes would be beneficial to the sender in this case. Yet, the following proposition shows that per-path monitoring (i.e., obtaining information regarding which paths successfully delivered the message) provides the sender with the same expected utility as per-hop monitoring. Thus, if the protocol provides per-path information, obtaining per-hop information yields no additional utility to the principal. This result is a simple integration of Proposition 3.1 and Claim 4.2.

*Proposition 4.4:* In parallel-paths topologies, the principal obtains the same expected utility under per-hop monitoring or per-path monitoring.

## V. Recursive Contracts

The focus so far was on *direct* contracts, where the principal contracts directly with each agent on the network and pays it directly. In *recursive* contracts, the principal contracts only with its direct neighbors (i.e., downstream routers), which in turn contract with their direct neighbors, and so on, all the way to the destination. Each agent contracts with its downstream agent once it receives the message, and makes the payment contingent on the final outcome. In this section, we study the effect of recursive contracts on the utility of the principal.

### A. Sequential Links

When the principal designs a contract for the first agent, he should take into account the incentives that the first agent needs to provide to the second agent, etc., all the way to

[4] For additional discussion on the appropriateness of different solution concepts in the context of online environments see, e.g., [9], [10].

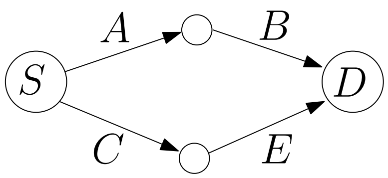

Fig. 3. Two parallel paths of length 2 each.

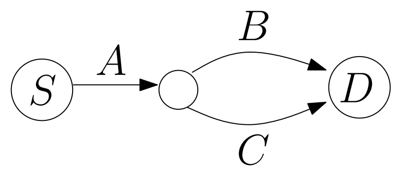

Fig. 4. A topology composed of a single link and two parallel links in sequence.

the destination. The next example demonstrates that in a sequential path of two links, the principal obtains the same expected utility under recursive and direct contracts.

*Example 5.1:* Consider a topology that is composed of two sequential links, $A$ and $B$. When the principal designs a contract with link $A$, he needs to consider the subsequent contract that $A$ should sign with $B$, which should satisfy the following constraint (here, $s_{A\to B}$ denotes the contract that A signs with B).

$$\textbf{IC} \quad E[s_{A\to B}|a_B = 1] - c \geq E[s_{A\to B}|a_B = 0]$$

where:

$$\begin{aligned} E[s_{A\to B}|a_B = 1] =& Pr(x^G_{B\to D}|a_B = 1)s^G_{A\to B} \\ &+ Pr(x^B_{B\to D}|a_B = 1)s^B_{A\to B} \end{aligned}$$

and

$$\begin{aligned} E[s_{A\to B}|a_B = 0] =& Pr(x^G_{B\to D}|a_B = 0)s^G_{A\to B} \\ &+ Pr(x^B_{B\to D}|a_B = 0)s^B_{A\to B} \end{aligned}$$

By substituting $s^B_i = 0$, we obtain the contract $s^G_{A\to B} = c/(1-k)$. Based on this payment, $S$ can express the IC constraint it should satisfy in a contract with $A$.

$$\begin{aligned} \textbf{IC} \quad & E[s_{S\to A} - s_{A\to B}|a_i = 1\forall i] - c \geq \\ & E[s_{S\to A} - s_{A\to B}|a_A = 0, a_B = 1] \end{aligned}$$

where:

$$\begin{aligned} & E[s_{S\to A} - s_{A\to B}|\forall i a_i = 1] = \\ & Pr(x^G_{A\to D}|\forall i a_i = 1)(s^G_{S\to A} - s^G_{A\to B}) + \\ & Pr(x^B_{A\to D}|a_i = 1\forall i)(s^B_{S\to A} - s^B_{A\to B}) \end{aligned}$$

and

$$\begin{aligned} & E[s_{S\to A} - s_{A\to B}|a_A = 0, a_B = 1] = \\ & Pr(x^G_{A\to D}|a_A = 0, a_B = 1)(s^G_{S\to A} - s^G_{A\to B}) + \\ & Pr(x^B_{A\to D}|a_A = 0, a_B = 1)(s^B_{S\to A} - s^B_{A\to B}) \end{aligned}$$

Solving for $s^G_{S\to A}$, we obtain: $s^G_{S\to A} = \frac{c(2-k)}{(1-k)^2}$. The expected total payment is

$$\begin{aligned} E[s] &= s^G_{S\to A} Pr(x^G_{S\to D}) = \frac{c(2-k)}{(1-k)^2}(1-k)^2 \\ &= c(2-k) \end{aligned} \tag{9}$$

Recall that in the direct case, the expected payment is (see Equation 8):

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

Substituting $n = 2$, we get:

$$E[s] = (1-k)^2 \left( \frac{c}{(1-k)^2} + \frac{c}{1-k} \right) = c(2-k)$$

which is equal to the expected total payment under recursive contracts. Since the expected benefit is equal as well, equal total payments also imply equal utility.

We next prove that this is a general result.

*Proposition 5.2:* In sequential-links scenarios, the expected principal's utility under the optimal direct contract is equal to that under the optimal recursive contract.

The intuition behind this result is that the agents in this scenario must provide appropriate incentives for their downstream agents, otherwise the message will not reach the destination with certainty.

Note that 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 link on the forwarding path and does not have to know the identities of the other links, it still needs to know the length of the path and the links' individual transit costs.

### B. *Other Network Topologies*

We next show that the result obtained for sequential-links topologies cannot be extended to general topologies. That is, a recursive design of the contract is more expensive for the principal than a direct contract. Consider, for example, the topology given in Figure 3.

*Claim 5.3:* For the parallel-paths topology given in figure 3, a recursive contract that induces all agents to exert effort is more expensive than a direct contract.

The next example demonstrates that for some network topologies, the relative performance of recursive contracts depend on the loss rate. Consider the network topology depicted in Figure 4, and suppose that the sender wishes to design a recursive contract that induces all three agents exert effort.

*Claim 5.4:* In the topology given in figure 4, there exists $k' \in (0,1)$, such that for any $k > k'$, the principal achieves the same expected utility under direct and recursive contracts, and for any $k \leq k'$, the principal achieves a lower expected utility under the best recursive contract.

## VI. QUALITY-OF-SERVICE (QOS) ROUTING

So far, we have considered the agents' strategy space to be limited to the drop and forward 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 network model where packets are forwarded on a best-effort or a priority basis [5]. 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 on the Internet. Individual autonomous systems (AS's) can either adopt hot-potato routing ($a = 0$), where a packet is handed off to the downstream AS at the earliest possible exit, or *cold-potato* routing ($a = 1$), where an AS carries the packet within its network as close to the destination as possible before handing it over to the downstream AS. In the absence of explicit incentives, it is not surprising that AS's choose hot-potato routing to minimize their costs, even though it may lead to suboptimal routes [25], [26].

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

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

$$Pr(x^G_{i\to i+1}|a_i) = 1 - (k - qa_i) \quad (10)$$

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

In what follows we show that in the QoS model, even in the sequential-links topology, the principal loses from his inability to observe individual actions.

*Proposition 6.1:* In the sequential-links topology, where agents decide whether to send traffic with high priority or low priority, the expected payment to each agent in the optimal contract exceeds its expected cost.

In section III, we found that in sequential-links topologies, the principal can achieve the optimal solution by observing only the final outcome if the actions of the agents are limited to drop versus forward. This is because the agents cannot free ride on others. If any single agent does not exert effort, a delivery failure is guaranteed and it obtains the low payment with certainty. In contrast, in the QoS model, as well as the parallel-links scenario, the agents have the opportunity to free ride on others, since the packet may reach the destination even if they do not exert effort. As a result, the principal cannot obtain the first best solution in the absence of monitoring. In these scenarios, per-hop monitoring improves the utility of the principal. In the absence of monitoring, the principal may achieve the optimal solution by imposing penalty on the agents upon failure (i.e., by setting $s^B_i < 0$), but the LL constraint eliminates this option. Note that other results that apply to the drop/forward model do not generalize to the QoS model. For example, even in the very basic topology of sequential links, the principal's utility under recursive contracts is lower than under direct ones.

## VII. Unknown Transit Costs

In certain network settings, the transit costs of links along the forwarding path may not be identical and may not be common knowledge, i.e., the transit costs may be hidden information. In this section, we study whether it is possible to design contracts that induce cooperative behavior in the presence of both hidden-action and hidden-information.

*Algorithmic mechanism design* addresses the problem of private information in computational settings by proposing mechanisms that induce rational agents to reveal their private information in computational settings. In their seminal paper [19], Nisan and Ronen presented a case study of routing in communication networks. In their setting, each router in the network has a private transit cost, and it acts strategically and declares a transit cost to maximize its profit. The goal is to design a mechanism that induces agents to reveal their true transit costs, such that the shortest path can be chosen. They propose a VCG-based mechanism [27] under which agents truthfully reveal their transit costs in dominant strategies. In [7], Feigenbaum et al. devised a distributed algorithm that computes the VCG payments. While the problem of hidden information is well handled, the problem of hidden action is not considered at all. They propose that the transit routers keep track of the amount of traffic routed through them via counters, and payments are periodically transferred from the principals to the transit links based on the counter values. It is simply assumed that transit links are obedient in packet forwarding behavior, and will not update the counters without exerting high effort in message forwarding.

In the remainder of this section, we ask whether it is possible for the principal to maximize social welfare given both hidden information and hidden action. More specifically, can a mechanism be designed such that (i) agents truthfully reveal their transit costs, and (ii) the contracted agants exert effort in forwarding the message.

If $k = 0$ (i.e., the network is lossless), the social welfare is the difference between the principal's benefit and the total costs incurred by the contracted agents . Thus, given a set $C$ of contracted agents, the social welfare is $W = w^G - \sum_{i\in C} c_i$. Clearly, choosing the shortest path maximizes social welfare. Therefore, a payment $p_i$ as proposed by [19] induces truthful revelation. In order to induce agents to engage in cooperative behavior as well, all that is needed is to make the payment contingent on arrival to the destination.

*Proposition 7.1:* In a lossless network, the contract $(s^G_i = p_i, s^B_i = 0)$, where $p_i$ is the VCG payment given in [19], induces a Nash equilibrium in which all the agents truthfully reveal their transit costs and the contracted agents exert effort in forwarding.

Like in the original problem (in environments of hidden information), maximizing the social welfare boils down to finding the shortest path; therefore, both the allocation and the payments can be computed in polynomial time. Yet, while the equilibrium achieved in settings with hidden information is a dominant-strategy equilibrium, the equilibrium achieved here (i.e., when the problems of hidden information and hidden actions are combined) is a Nash equilibrium, which is a weaker solution concept. The problem of inducing agents to truthfully reveal their private information and exert effort in forwarding under lossy networks (i.e., where $k > 0$) is still an open problem.

## VIII. Case Study: Internet Routing

Inter-domain routing over the Internet involves the forwarding of packets over multiple network domains or Autonomous systems (AS's). This implies some form of economic incentives is needed to induce one network to carry traffic for another network.

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) is faced with a scenario characterized by both hidden action and hidden information. Such contracts involve paying some premium over the real cost, and it is not clear whether recursive contacts 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 [31]. Yet another approach is to delegate the task of obtaining topology and performance information to a third-party routing service provider [15]. These proposals are part of a broader approach to provide end hosts with greater control over the packets they send. 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 setting in sections III and IV. If the principal sends packets over a multi-hop path, corresponding to the sequential-links topology, then optimal contracts can be implemented with direct contracts with the individual agents along the path. Moreover, as we have shown in Section V, the optimal solution can also be implemented with recursive contracts as long as each agent can choose the next hop. In either case, monitoring information cannot further improve the principal's utility. This suggests that implementing source routing is not necessary in these architectures. Not providing source routing would lower the complexity of the architectures and eliminate the security concerns associated with source routing.

While the various proposals for acquiring network topology and performance information might imply that we can treat the network information as common knowledge and apply the mechanisms proposed in the early sections of the paper, the approach taken by these proposals is different than the mechanism design approach in that they 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 [2]. Another approach is to have third parties like Keynote independently monitor the network performance.

## IX. Conclusions

Recent advances in the field of algorithmic mechanism design addresses the issue of private information held by selfish agents in computational settings in general and in routing scenarios in particular. In this paper we tackle the complementary problem, that of hidden actions performed by the intermediate routers in network routing. This is a first step toward a general analysis of hidden actions in computational settings. To deal with hidden actions in such settings, we employ the principal-agent model with multiple agents. The solution is based on the observation that a properly designed contract, in which the payments are contingent upon the final outcome, can influence a rational agent to exert effort in forwarding a packet on behalf of the sender. We present the structure of the optimal contracts under various network topologies, and show that information about the intermediate outcomes within the network (i.e., monitoring information) may or may not improve the utility of the sender, depending on the network topology. In addition, we show that the contracts can be implemented either directly with each router on the path, or recursively such that each router contracts only with its downstream router. Depending on the network topology, the recursive implementation may or may not result in a loss in the principal's utility. Finally, we study the design of contracts in scenarios that are characterized by both hidden actions and hidden information, and show that it is possible to induce an equilibrium in which all agents truthfully reveal their private information and engage in the desired actions.

## X. Acknowledgements

This work is supported in part by the National Science Foundation under ITR awards ANI-0085879 and ANI-0331659, and Career award ANI-0133811.

## Appendix

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

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

*the expected principal's utility in the longer path is higher than his expected utility in the shorter path.*

*Proof:* Since the probability of success is equal in the two paths, it is sufficient to show that the total payments are the same. The total payment in path $j$ is given by: The expected payment in path $j$ is given by:

$$\sum_{i=1}^{n_j} s_i^G = \sum_{i=1}^{n_j} \frac{c_j}{(1-k_j)^{n_j-i+1}}$$

We get:

$$\sum_{i=1}^{n_1} s_i^G > \sum_{i=1}^{n_2} s_i^G$$

$$\Updownarrow$$

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

$$\Updownarrow$$

$$c_1 \frac{1-(1-k_1)^{n_1}}{k_1} - c_2 \frac{1-(1-k_2)^{n_2}}{k_2} > 0$$

$$\Updownarrow$$

$$\frac{c_1 k_2 (1-(1-k_1)^{n_1}) - c_2 k_1 (1-(1-k_2)^{n_2})}{k_1 k_2} > 0$$

By Equation 2, the above inequation holds if and only if

$$c_1 k_2 - c_2 k_1 > 0$$

which (by Equation 1) holds if and only if

$$n_2 k_2 - n_1 k_1 > 0$$

By equation 2, $n_1 = n_2 \frac{\ln(1-k_2)}{\ln(1-k_1)}$. Substituting for $n_1$, we get:

$$n_2 k_2 - n_1 k_1 > 0$$

$$\Updownarrow$$

$$n_2 k_2 - n_2 \frac{\ln(1-k_2)}{\ln(1-k_1)} > 0$$

$$\Updownarrow$$

$$\frac{k_2 \ln(1-k_1) - k_1 \ln(1-k_2)}{\ln(1-k_1)} > 0$$

$$\Updownarrow$$

$$k_2 \ln(1-k_1) - k_1 \ln(1-k_2) < 0$$

But from Equations 2 and 3, it follows that $k_1 > k_2$. Therefore, $\ln(1-k1) < \ln(1-k_2)$, and therefore $k_2 \ln(1-k_1) - k_1 \ln(1-k_2) < 0$, as required. ■

**Proposition 3.3** *In the parallel links scenario, under the optimal contract that induces high-effort behavior from all intermediate edges in the Nash Equilibrium, the expected payment to each node exceeds its expected cost if the limited-liability constraint is imposed, and equals to its expected cost otherwise.*

*Proof:* We prove the following:

- If the limited liability (LL) constraint is not imposed, then the optimal payment schedule is:

$$s_i^B = -\frac{c(k^{1-n}-1)}{1-k} \quad ; \quad s_i^G = \frac{c}{1-k}$$

  and the expected payment to each agent equals its expected cost.
- If the LL constraint is imposed, the optimal payment schedule is:

$$s_i^B = 0 \quad ; \quad s_i^G = \frac{c}{k^{n-1}(1-k)}$$

  and the expected payment to each agent exceeds its expected cost.

When the LL constraint is not imposed, the $IC$ and $IR$ constraints for each agent $i$ can be expressed as follows:

$$\textbf{IC} \quad Pr(x^G|\forall j a_j = 1)s_i^G + Pr(x^B|\forall j a_j = 1)s_i^B - c \geq Pr(x^G|a_i = 0, a_{j\neq i} = 1)s_i^G + Pr(x^B|a_i = 0, a_{j\neq i} = 1)s_i^B \quad (11)$$

$$\textbf{IR} \quad Pr(x^G|\forall j a_j = 1)s_i^G + Pr(x^B|\forall j a_j = 1)s_i^B - c \geq 0 \quad (12)$$

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

$$\textbf{IC} \quad (1-k^n)s_i^G + k^n s_i^B - c \geq (1-k^{n-1})S_i^G + k^{n-1}s_i^B$$

$$\textbf{IR} \quad (1-k^n)s_i^G + k^n s_i^B - c \geq 0$$

Both constraints bind, and by solving the two equations, we obtain the payments that are given above. The expected cost of each link is $c$, since any link is a direct neighbor of the sender and thus gets the message with certainty. The expected payment to agent $i$ is given by:

$$E[s_i] = (1-k^n)s_i^G + k^n s_i^B = c$$

Thus, if the LL constraint is not imposed, the principal can achieve a contract in which he pays each link exactly its expected payment. Note that the payment upon a delivery failure is negative in this case (i.e., the agents pay the sender).

If the LL constraint is imposed, then in the optimal contract $s_i^B = 0$. It is easy to verify that the IC constraint binds. Substituting $s_i^B = 0$ into the IC constraint, we obtain the payment $s_i^G = \frac{c}{k^{n-1}(1-k)}$. The expected cost of each agent $i$ is $c$, while its expected payment is:

$$E[s_i] = c\frac{1-k^n}{k^{n-1}(1-k)}$$

To show that $E[s_i] > c$, it is sufficient to show that $\frac{1-k^n}{k^{n-1}(1-k)} > 1$.

$$\frac{1-k^n}{k^{n-1}(1-k)} > \Leftrightarrow 1-k^n > k^{n-1}(1-k) \Leftrightarrow 1-k^{n-1} > 0$$

which holds for any $k < 1$. Thus, $E[s_i] > E[c_i]$ for any $k < 1$. ■

**Proposition 4.1** *In the sequential-links topology, 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 incented 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 the first-best solution. ■

**Claim 4.2** *In sequential-links topologies, a mechanism that makes payments contingent on arrival to the next link (i.e., pays $s_i^G$ if the message reached link $i+1$, and pays $s_i^B$ otherwise) yields the same expected utility to the principal as one that makes the payments contingent on arrival to the destination.*

*Proof:* 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. If the payments are contingent on individual success, the principal needs to satisfy the following constraint:

$$\textbf{IC} \quad \begin{aligned} &Pr(x_i^G|a_i=1)s_i^G + Pr(x_i^B|a_i=1)s_i^B - c \geq \\ &Pr(x_i^G|a_i=0)s_i^G + Pr(x_i^B|a_i=0)s_i^B \end{aligned} \tag{13}$$

Substituting for the probabilities, we can express the constraint as:

$$\textbf{IC} \quad (1-k)s_i^G + ks_i^B - c \geq s_i^B$$

which yield the optimal contract of $s_i^B = 0$ and $s_i^G = \frac{c}{1-k}$.

The expected total payment under this scenario is:

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

as in the scenario without monitoring (see Equation 8.) ■

**Proposition 4.3** *In any network topology, when per-hop monitoring is available, the principal achieves a contract in which he pays each agent his expected cost. The payment schedule is $s_i^B = 0$ and $s_i^G = \frac{c}{1-k}$.*

*Proof:* The IC constraint can be expressed as:

$$\textbf{IC} \quad (1-k)s_i^G + ks_i^B - c \geq s_i^B$$

which yield the payments $s_i^B = 0$ and $s_i^G = \frac{c}{1-k}$, as required. ■

**Proposition 4.4** *In parallel-paths topologies, the principal obtains the same expected utility under per-hop monitoring or per-path monitoring.*

*Proof:* We first consider the case of per-hop monitoring. Under per-hop monitoring, the principal needs to satisfy the $IC$ and $IR$ constraints given in Equations 6 and 7, which yields the optimal payments $s_i^G = \frac{c}{1-k}$, $S_i^B = 0$. Under these payments, as shown in Section III, the expected payment to each agent equals its expected cost.

In the case of per-path monitoring, the principal needs to satisfy the $IC$ and $IR$ constraints given in Equations 4 and 5, which yield the payments $S_i^B = 0$ and $S_i^G = \frac{c}{(1-k)^{n-i+1}}$. Under these payments, as shown in Section III, the expected payment to each agent equals its expected cost as well. ■

**Proposition 5.2** *In sequential-links scenarios, the expected principal's utility under the optimal direct contract is equal to that under the optimal recursive contract.*

*Proof:* Since the principal's expected benefit in both cases is the probability of success given that all the agents exert effort multiplied by the principal's value from a successful transmission, we only have to show that the expected total payments are equal. In a recursive contract, each agent contracts with its downstream agent once it receives the message. Thus, it incurs the cost with certainty and gains the difference between its payment and the payment it needs to transfer if the message reaches the destination. In a contract with agent $i$ (the $i$'th agent on the path), the following constraint should be satisfied:

$$Pr(x_{i\to D}^G|\forall j a_j = 1)(s_i^G - s_{i+1}^G) - c_i = 0$$

substituting the probability of success, we get for all $1 \leq i \leq n$:

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

Substituting these values recursively, we get:

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

and the expected total payment is

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

which is equal to the total expected payment in direct payments, as expressed in Eq. 8. ■

**Claim 5.3** *For the parallel-paths topology given in figure 3, a recursive contract that induces all agents to exert effort is more expensive than a direct contract.*

*Proof:* Consider the constraint that agent $A$ needs to satisfy in order to induce agent $B$ to exert effort. Here, $\hat{s}$ denotes the payment in the recursive case.

$$\left(1 - \left(1-(1-k)^2\right)k\right)\hat{s}_{A\to B}^G - c \geq (1-k)^2 \hat{s}_{A\to B}^G$$

which yields: $s_{A\to B}^G = \frac{c}{(1-k)(1-(1-k)^2)}$. Given this payment, the constraint that the principal needs to satisfy in order to induce $A$ to both exert effort and induce $B$ to exert effort is given by:

$$\begin{aligned} &\left(1 - \left(1-(1-k)^2\right)^2\right)\left(\hat{s}_{S\to A}^G - \hat{s}_{A\to B}^G\right) - c \geq \\ &(1-k)^2 \hat{s}_{S\to A}^G \end{aligned}$$

which yields:

$$\begin{aligned} \hat{s}_{S\to A}^G &= \frac{c}{(1-k)^2(1-(1-k)^2)} \\ &+ \frac{c}{(1-k)(1-(1-k)^2)} \cdot \frac{2-(1-k)^2}{1-(1-k)^2} \end{aligned}$$

The recursive contract with $C$ is symmetric, therefore, the total payment in case of a successful transmission is $2 \cdot s^G_{S\to A}$.

Now consider the total payment under a direct contract. In order to induce agent $A$ to exert effort, the principal needs to satisfy the following constraint:

$$\left(1-\left(1-(1-k)^2\right)^2\right) s^G_{S\to A} - c \geq (1-k)^2 s^G_{S\to A}$$

which yields: $s^G_{S\to A} = \frac{c}{(1-k)^2(1-(1-k)^2)}$. Similarly, in order to induce $B$ to exert effort, he needs to satisfy the following constraint:

$$\left(1-\left(1-(1-k)^2\right)k\right) s^G_{S\to B} - c \geq (1-k)^2 s^G_{S\to B}$$

which yields: $s^G_{S\to B} = \frac{c}{(1-k)(1-(1-k)^2)}$. Therefore, the total payment in case of a successful transmission is $2 \cdot \left(s^G_{S\to A} + s^G_{S\to B}\right)$. But since $\frac{2-(1-k)^2}{1-(1-k)^2} > 1$ for any $k$, the total payment under recursive contracts is greater than under direct contracts. ■

**Claim 5.4** *In the topology given in figure 4, there exists $k' \in (0,1)$, such that for any $k > k'$, the principal achieves the same expected utility under direct and recursive contracts, and for any $k \leq k'$, the principal achieves a lower expected utility under the best recursive contract.*

*Proof:* Consider first the total payment under success in direct contracts. The payment to $A$ is derived through the constraint: $(1-k)(1-k^2)s^G_A - c \geq 0$, yielding $s^G_A = \frac{c}{(1-k)(1-k^2)}$. The payment to agent $B$ (and similarly $C$) is derived through the constraint: $(1-k^2)s^G_B - c \geq (1-k)s^G_B$, yielding $s^G_B = \frac{c}{k(1-k)}$. Thus, the total payment is $s^G_A + 2\cdot s^G_B = \frac{c}{(1-k)(1-k^2)} + 2 \cdot \frac{c}{k(1-k)}$.

Now consider the optimal recursive contract that induces all agents to exert effort. If $A$ does not contract with any of the agents, or if he, himself, does not exert effort, then his utility is 0.

If $A$ contracts with a single agent (say $B$), he needs to satisfy $(1-k)\hat{s}^G_{A\to B} - c \geq 0$, yielding $\hat{s}^G_{A\to B} = \frac{c}{1-k}$. Therefore, $A$'s expected utility if he contracts with a single agent is: $E[u_A] = (1-k)(1-k)\left(\hat{s}^G_{S\to A} - \frac{c}{1-k}\right) - c$.

If $A$ induces both agents to exert effort, the payment to agent $B$ (and similarly $C$) is derived through the constraint: $(1-k^2)s^G_B - c \geq (1-k)s^G_B$, yielding $s^G_B = \frac{c}{k(1-k)}$. Therefore, $A$'s expected utility if he contracts with both agents is: $E[u'_A] = (1-k)(1-k^2)\left(\hat{s}^G_{S\to A} - 2\cdot\frac{c}{k(1-k)}\right) - c$.

Therefore, in order for the principal to induce $A$ to exert effort and induce both agents to exert effort, he needs to satisfy: $E[u'_A] \geq E[u_A]$, yielding: $\hat{s}^G_{S\to A} = \frac{2c(1+k)}{k^2(1-k)} - \frac{c}{k(1-k)}$, and also $E[u'_A] \geq 0$, yielding: $\hat{s}^G_{S\to A} = \frac{c}{(1-k)(1-k^2)} + \frac{2c}{k(1-k)}$ (which is equal to the total payment in the direct case).

We conclude that for any $k$, the total payments under recursive contracts is at least that under direct contracts, but it may also be higher. This happens when:

$$\frac{2c(1+k)}{k^2(1-k)} - \frac{c}{k(1-k)} > \frac{c}{(1-k)(1-k^2)} + \frac{2c}{k(1-k)}$$

$$\Updownarrow$$

$$k^4 - 2k^3 - 6k^2 + 2 > 0$$

Let $f(k) = k^4 - 2k^3 - 6k^2 + 2$. It is easy to verify that its derivative is negative for any $k \in [0,1]$, therefore $f(k)$ is monotonically decreasing in $k$. In addition, $f(0) > 0$ while $f(1) < 0$. Thus, by the intermediate value theorem, there exists $k'$ for which $f(k) = 0$, for any $k \leq k'$, $f(k) \geq 0$ (i.e., the total payment under recursive contracts is higher), and for any $k > k'$, $f(k) < 0$, thus recursive contracts result in the same payments as direct ones. ■

**Proposition 6.1** *In the sequential-links topology, where agents decide whether to send traffic with high priority or low priority, the expected payment to each agent in the optimal contract exceeds its expected cost.*

*Proof:*

The IC constraint is the same as specified in the proof of proposition 3.1, but the probabilities change, based on Eq. 10, as follows:

$$\begin{aligned}\textbf{IC}\quad & (1-k+q)^{n-i+1}s^G_i + \\ & (1-(1-k+q)^{n-i+1})s^B_i - c \geq \\ & (1-k)(1-k+q)^{n-i}s^G_i + \\ & (1-(1-k)(1-k+q)^{n-i})s^B_i\end{aligned}$$

This constraint yields the payment $s^G_i = \frac{c}{(1-k+q)^{n-i}q}$, and the expected payment is thus given by:

$$E[s_i] = (1-k+q)^n s^G_i = (1-k+q)^i \frac{c}{q}$$

While the expected cost is given by:

$$E[c_i] = (1-k+q)^{i-1}c$$

To show that the expected payment exceeds the expected cost, we need to show that $E[s_i] - E[c_i] > 0$. That is,

$$(1-k+q)^i \frac{c}{q} - (1-k+q)^{i-1}c > 0$$

$$\Updownarrow$$

$$c(1-k+q)^{i-1}\left(\frac{1-k+q}{q} - 1\right) > 0$$

$$\Updownarrow$$

$$1-k > 0$$

which holds for any $k \in (0,1)$. ■

**Proposition 7.1** *In a lossless network, the contract $(s^G_i = p_i, s^B_i = 0)$, where $p_i$ is the VCG payment given in [19], induces a Nash equilibrium in which all the agents truthfully reveal their transit costs and the contracted agents exert effort in forwarding.*

*Proof:* We have to show that no agent can increase his utility by (i) truthfully revealing his cost, but dropping the message, (ii) misreporting his cost, and dropping the message, or (iii) misreporting his cost, but forwarding the message. The equilibrium utility of an agent who truthfully reveals his cost and forwards the message is $p_i - c_i$ which is a non-negative value, based on regular VCG reasoning (see [19]). If an agent engages in misbehaviors (i) or (ii), his utility becomes 0, and he forgoes a non-negative utility. If, on the other hand, he forwards the message, then he is best-off revealing his true cost due to regular VCG reasoning. Thus, misbehavior (iii) cannot improve his utility either. ■

## REFERENCES


[1] Andersen, D. G., Balakrishnan, H., Kaashoek, M. F., and Morris, R. Resilient Overlay Networks, in *18th ACM SOSP* (2001).

[2] Argyraki, K., Maniatis, P., Cheriton, D., and Shenker, S. Providing Packet Obituaries, in *Third Workshop on Hot Topics in Networks (HotNets)* (November 2004).

[3] Babaioff, M., Feldman, M., and Nisan, N. Combinatorial Agency, in *ACM Conference on Electronic Commerce (EC'06)* (2006).

[4] Bansal, S., and Baker, M. Observation-based cooperation enforcement in ad hoc networks. *Technical report, Stanford university* (2003).

[5] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and Weiss, W. An Architecture for Differentiated Service. RFC 2475, 1998.

[6] Buchegger, S., and LeBoudec, J.-Y. Coping with False Accusations in Misbehavior Reputation Systems For Mobile Ad Hoc Networks, in *EPFL, Technical report* (2003).

[7] 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).

[8] Feigenbaum, J., Sami, R., and Shenker, S. Mechanism Design for Policy Routing, in *Yale University, Technical Report* (2003).

[9] 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).

[10] Friedman, E., and Shenker, S. Learning and implementation on the internet, in *Manuscript. New Brunswick: Rutgers University, Department of Economics* (1997).

[11] Ganesan, D., Govindan, R., Shenker, S., and Estrin, D. Highly-resilient, energy-efficient multipath routing in wireless sensor networks. *SIGMOBILE Mob. Comput. Commun. Rev. 5*, 4 (2001), 11–25.

[12] Goldschlag, D., Reed, M., and Syverson, P. Onion Routing for Anonymous and Private Internet Connections, in *Communications of the ACM* (1999), vol. 42.

[13] Holmstrom, B. Moral Hazard in Teams. *Bell Journal of Economics 13* (1982), 324–340.

[14] Jakobsson, M., Hubaux, J.-P., and Buttyan, L. A Micro-Payment Scheme Encouraging Collaboration in Multi-Hop Cellular Networks, in *Financial Cryptography* (2003).

[15] Lakshminarayanan, K., Stoica, I., and Shenker, S. Routing as a service, in *UCB Technical Report No. UCB/CSD-04-1327* (January 2004).

[16] Legeros, P., and Matthews, S. Efficient and Nearly-Efficient Partnerships. *The Review of Economic Studies 60*, 3 (1993), 599–611.

[17] Mass-Colell, A., Whinston, M., and Green, J. *Microeconomic Theory*. Oxford University Press, 1995.

[18] Maxemchuk, M. Dispersity routing in highspeed networks. *Computer Networks and ISDN Systems 25* (1993), 645–661.

[19] Nisan, N., and Ronen, A. Algorithmic Mechanism Design, in *Proceedings of the 31st Symposium on Theory of Computing* (1999).

[20] Raghavan, B., and Snoeren, A. Priority Forwarding in Ad Hoc Networks with Self-Ineterested Parties, in *Workshop on Economics of Peer-to-Peer Systems* (June 2003).

[21] Rasmusen, E. Moral Hazard in Risk-Averse Teams. *Rand Journal of Economics 18* (1987), 428–435.

[22] Ratnasamy, S., Francis, P., Handley, M., Karp, R., and Shenker, S. A Scalable Content-Addressable Network, in *Proceedings of ACM SIGCOM* (September 2001).

[23] Stoica, I., Morris, R., Karger, D., Kaashoek, M. F., and Balakrishnan, H. Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications, in *Proceedings of ACM SIGCOM* (September 2001).

[24] strausz, R. Moral Hazard in Sequential Teams. *Departmental Working Paper. Free University of Berlin* (1996).

[25] Teixeira, R., Griffin, T., Shaikh, A., and Voelker, G. Network sensitivity to hot-potato disruptions, in *Proceedings of ACM SIGCOMM* (September 2004).

[26] Teixeira, R., Shaikh, A., Griffin, T., and Rexford, J. Dynamics of hot-potato routing in IP networks, in *Proceedings of ACM SIGMETRICS* (June 2004).

[27] Vickrey, W. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* (1961), 8–37.

[28] Wang, W., and Li, X.-Y. Low-Cost Routing in Selfish and Rational Wireless Ad-Hoc Networks. *IEEE Trans. Mobile Computing 5*, 5 (2006), 596–607.

[29] Winter, E. Incentives and Discrimination. *American Economic Review 94* (2004), 764–773.

[30] Zhong, S., Li, L. E., Yang, Y. R., and Liu, Y. On Designing Incentive-Compatible Routing and Forwarding Protocols in Wireless Ad-Hoc Networks – An Integrated Approach Using Game Theoretical and Cryptographic Techniques. *ACM/Baltzer Wireless Networks (WINET) Journal 5*, 5 (2006), 596–607.

[31] Zhu, D., Gritter, M., and Cheriton, D. Feedback-based Routing, in *Proc Hotnets-I* (2002).


**Michal Feldman** is a postdoc fellow at the School of Computer Science and Engineering at the Hebrew University of Jerusalem and at Tel-Aviv University. She received her PhD from the School of Information Management and Systems at the University of California at Berkeley in 2005. Her research interests lie in the intersection of Game Theory, Microeconomic Theory and Communication Networks. She is the recipient of the Lady Davis Fellowship (2005-2007).

**John Chuang** is Associate Professor of Information Management and Systems at the University of California, Berkeley. He received a B.S. and M.S. in Electrical Engineering from the University at Southern California and Stanford University respectively, and a Ph.D. in Engineering and Public Policy from Carnegie Mellon University. His research focus is on economics-informed design of computer networks and distributed systems.

**Ion Stoica** received his PhD from the Carnegie Mellon University in 2000. He is an Assistant Professor in the EECS Department at University of California at Berkeley, where he does research on peer-to-peer network technologies in the Internet, resource management, and network architectures. Stoica is the recipient of a Sloan Foundation Fellowship (2003), a Presidential Early Career Award for Scientists & Engineers (PECASE) (2002), and of the ACM doctoral dissertation award (2001). He is a member of ACM.real time systems.

**Scott Shenker** spent his academic youth studying theoretical physics but soon gave up chaos theory for computer science. Continuing to display a remarkably short attention span, his research over the years has wandered from computer performance modeling and computer networks to game theory and economics. Unable to hold a steady job, he currently splits his time between the U. C. Berkeley Computer Science Department and the International Computer Science Institute (ICSI).