Hidden-Action in Multi-Hop Routing

Michal Feldman John Chuang

School of Information Management and Systems
University of California, Berkeley

Abstract – In any multi-hop routing scheme, cooperation by the intermediate nodes are essential for the succesful delivery of traffic. However, the effort exerted by the intermediate nodes are often unobservable by the source and/or destination nodes. We show it is possible to overcome this problem of *hidden action* by designing contracts, in the form of payments, to induce cooperation from the intermediate nodes. Interestingly, the ability to monitor per-hop or per-path outcomes, even if costless to implement, may not improve the welfare of the participants or the performance of the network.

1 Introduction

In a multi-hop network, a sender relies on intermediate nodes on a routing path to forward packets toward the destination. The intermediate nodes may incur significant communication and computation costs in the forwarding of packets without deriving any direct benefit from doing so [9]. 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.

The sender can provide incentives, e.g., in the form of payments, to encourage the intermediate nodes to forward its packets. However, the actions of the intermediate nodes are often *hidden* from the sender. In many cases, the sender 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 the sender to pinpoint the location of the failure, the sender 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.

This problem of hidden action arises in a variety of peer-to-peer network contexts. In application-layer overlay networks for resilient routing [1], anonymous routing [14], as well as those implementing distributed hash tables (DHTs) [23, 24, 27], intermediate nodes may choose not to forward traffic for others. In wireless ad hoc networks, nodes may strategically drop packets to conserve their constrained energy resources. Monitoring is difficult to accomplish given the large scale and the dynamic memberships and topologies of these networks. Even in Internet routing, network routers may drop packets for a variety of legitimate reasons without notification of the senders, and end users have only crude forms of monitoring at their disposal. Indeed, the problem of hidden action arises in many non-networking-related contexts as well. Also known as *moral hazard*, the problem is well studied in the economics literature on information asymmetry and agency theory.

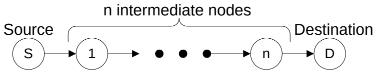

Figure 1: Multi-hop path from sender to destination.

By adopting the *principal-agent* model, with multiple agents making sequential hidden actions [15, 28], we can study the problem and possible resolution of hidden action in multi-hop routing. 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 in different priority levels. Second, we find that the absence of monitoring may have no impact on the sender's expected utility. For a sender 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. Yet, the use of per-hop monitoring makes the system more vulnerable to several forms of collusion among the intermediate nodes. Finally, we consider network scenarios where multiple disjoint paths exist between the sender and the receiver. If the sender elects to send multiple copies of the packets to maximize the likelihood of delivery, it will incur the same expected costs whether it obtains path-specific monitoring information or not.

2 Model

The scenarios we consider in this work consist of a single sender (the *principal*) that wishes to transmit traffic to a destination along a path that is composed of n intermediate transit nodes (the *agents*), as shown in figure 1[1]. The sender makes individual *take-it-or-leave-it* offers to the intermediate nodes, and the nodes choose between low and high effort actions sequentially, based on the payment schedule, to maximize their expected payoff.

Each agent chooses an action, $a \in \{0, 1\}$, where $a = 0$ and $a = 1$ stand for the low-effort and high-effort actions, respectively. Each action is associated with a cost $C(a)$, such that:

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

We denote $C(a = 1) = c$,and normalize the cost of the low-

---

[1] In the remainder of the paper, we use the terms principal and sender interchangeably, as well as the terms agents and intermediate nodes.

effort action to zero [2].

The outcome is denoted by $x \in \{x^H, x^L\}$, where x^H and x^L stand for the respective outcomes in which the packet reaches or does not reach the destination [3]. The benefit of the sender from the outcome is denoted by $b(x)$, where:

$$b(x^H) = b^H$$

$$b(x^L) = b^L = 0$$

We assume that the principal cannot observe the actions of the agents, thus can only make the payment contingent on the outcome.

The utility of agent i, denoted by u_i, is the difference between the payment it receives (s_i) and the cost it incurs from performing the action ($C(a_i)$):

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

The utility of the principal, denoted by w, is the difference between the benefit it derives from the outcome ($b(x)$) and the sum of the payments to the entire agent population:

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

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

Finally, we make the following simplifying assumptions:

- Transit costs are common knowledge and homogeneous. While the assumption about common knowledge is essential to our model, the homogeneity assumption is for presentation clarity only, and the model can easily account for heterogeneous costs.
- The number of nodes along the path and their order is common knowledge.
- Nodes are risk-neutral.
- The $n+1$ per-hop transmission events from the source to the destination are independent and identically distributed (i.i.d.).

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): the expected utility from participation should (weakly) exceed its reservation utility (which is normalized to 0).

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

3 Results

In this section we analyze various scenarios. Scenario $v1$ refers to cases in which agents decide whether to drop or forward packets. This scenario is studied in sections 3.1 and 3.2, where $v1a$ and $v1b$ refer to cases without and with per-hop monitoring, respectively. In section 3.3 we consider scenarios in which agents decide whether to engage in priority or best-effort forwarding ($v2$). Finally, in section 3.4, we consider scenarios in which multiple disjoint paths exist from the source to the destination ($v3$).

3.1 Drop versus Forward without Monitoring ($v1a$)

In this scenario the agents decide whether to drop ($a = 0$) or forward ($a = 1$) packets. The loss rate in each hop is k. Thus, the probability that a packet successfully gets from node i to its successor, $i+1$, can be expressed as:

$$Pr(x^H_{i \to i+1} | a_i) = (1-k)a_i \tag{1}$$

where $x^H_{i \to j}$ denotes a successful transmission from node i to node j.

The sender can only observe the final outcome, that is, whether or not the packet has reached the destination. Consequently, the sender makes the payment schedule to each agent contingent on the final outcome, x, as follows:

$$s_i(x) = (s_i^L, s_i^H)$$

where:

$$s_i^L = s_i(x = x^L)$$
$$s_i^H = s_i(x = x^H)$$

Proposition 3.1 *Under the best contract that induces high-effort behavior from all intermediate nodes in the Nash Equilibrium:*

- *The expected payment to each node equals its expected cost.*
- *The payment schedule is:*

$$s_i^L = 0 \tag{2}$$

$$s_i^H = \frac{c}{(1-k)^{n-i+1}} \tag{3}$$

- *The principal obtains the same expected utility as under the first-best contract (where actions are observed).*

Proof The sender needs to satisfy the IC and IR constraints for each agent i, which can be expressed as follows:

$$(IC): Pr(x^H | a_{j \geq i} = 1)s_i^H + Pr(x^L | a_{j \geq i} = 1)s_i^L - c \geq$$
$$Pr(x^H | a_i = 0, a_{j > i} = 1)s_i^H + Pr(x^L | a_i = 0, a_{j > i} = 1)s_i^L$$

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.

---

[2] The main results that we obtain in this work are not affected by this normalization.

[3] We use the terminology of "packet" to refer to the traffic transmitted, but depending on the application, it can refer to any traffic unit.

$(IR): Pr(x^H_{S\to i}|a_{j<i}=1)(Pr(x^H|a_{j\geq i}=1)s_i^H +$
$Pr(x^L|a_{j\geq i}=1)s_i^L - c) + Pr(x^L_{S\to i}|a_{j<i}=1)s_i^L \geq 0$

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

Based on equation 1, The above constraints can be expressed as follows:

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

It is a standard result that both constraints bind at the best contract [20]. Solving the two equations, we obtain the solution that is presented in equations 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:

$$\begin{aligned} E[s_i] &= Pr(x^H)s_i^H + Pr(x^L)s_i^L \\ &= (1-k)^{n+1}\frac{c}{(1-k)^{n-i+1}} \\ &= (1-k)^i c \end{aligned}$$

Finally, it is easy to see that the principal obtains the same expected utility as under full information. Were the actions observable, the sender would pay each node that receives the packet its transit cost, c, for forwarding the packet, thus the expected payment to each node would be $(1-k)^i c$. ■

The sender will enter into a contract only if the expected benefit it derives from a successful outcome exceeds its expected implementation cost. Thus, there is a threshold value of b^H, denoted by b_t^H, below which the sender will not enter the contract. This value increases exponentially in the number of nodes, n, and polynomially with the loss rate, k (see Appendix A). Intuitively, the longer the path and the higher the loss rate, the higher the implementation cost will become.

Based on these results, we can already make several interesting observations:

- The sequential nature of the transmission reduces the sender's implementation cost, because the sender needs to compensate each node only for its *expected* transit cost, which is the product of its transit cost, c, and the probability that the packet reaches that node (< 1).
- This Nash equilibrium has nice properties with respect to vulnerability to *collusion*. In particular, two or more intermediate nodes cannot obtain better payoffs by colluding with each other, since nodes receive the high payment only if the packet reaches the destination.

3.2 Drop versus Forward with Monitoring ($v1b$)

In this section, we consider the scenario in which the sender obtains per-hop monitoring information. That is, as opposed to $v1a$, where the sender could only observe the final outcome, under $v1b$, the sender is notified about every node along the path that receives the packet. In what follows, we assume that employing per-hop monitoring information is costless. While this assumption may not be realistic, it provides us with the best case scenario for the use of monitoring.

As before, the sender wishes to induce an equilibrium in which all intermediate nodes forward the packet. This scenario essentially reduces to n identical instances of a single principal - single agent problem. It is a well-known result (see e.g., [20]) that if the principal and the agent are risk-neutral, the expected payment paid by the principal equals the payment under full information.

Proposition 3.2 *The sender derives the same expected utility in equilibrium whether it obtains per-hop monitoring information or not.*

Proof The expected utility of the sender 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 sender has to satisfy the IC and IR constraints. In order to avoid confusion between the scenarios $v1a$ and $v1b$, we denote the payment to node i upon successful delivery to the next hop by s_i^1, and that upon failure by s_i^0.

$(IC): Pr(x^H_{i\to i+1}|a_i=1)s_i^1 + Pr(x^L_{i\to i+1}|a_i=1)s_i^0 - c \geq$
$Pr(x^H_{i\to i+1}|a_i=0)s_i^1 + Pr(x^L_{i\to i+1}|a_i=0)s_i^0$
$(IR): Pr(x^H_{S\to i}|a_{j<i}=1)(Pr(x^H_{i\to i+1}|a_i=1)s_i^1$
$+Pr(x^L_{i\to i+1}|a_i=1)s_i^0 - c) \geq 0$

These constraints can be expressed as follows:

$(IC): (1-k)s_i^1 + ks_i^0 - c \geq s_i^0$
$(IR): (1-k)^i((1-k)s_i^1 + ks_i^0 - c) \geq 0$

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

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

The expected total payment under this scenario is:

$$E[S] = \sum_{i=1}^{n}((1-k)^i k(s^0 + (i-1)s^1)) + (1-k)^{n+1}ns^1$$

which yields: $\frac{c(1-k-(1-k)^n+(1-k)^n k)}{k}$, as in $v1a$. ■

In both scenarios, the expected payment to each node equals its expected cost, $(1-k)^i c$. Hence, the sender incurs a total cost of $c\sum_{i=1}^{n} Pr(x^H_{S\to i})$. To see this in $v1b$, recall that a node faces the payment schedule only if it gets the packet. Therefore, we just need to show that the expected payment to a node that received the packet, denoted by $E[s_i|x^H_{S\to i}]$, equals c.

$E[s_i|x^H_{S\to i}] = Pr(x^H_{i\to i+1}|a_i=1)s_i^1 + Pr(x^L_{i\to i+1}|a_i=1)s_i^0$

$$= (1-k)\frac{c}{1-k} = c, \text{ as required.}$$

Many of the works in the field of incentive mechanisms for multi-hop routing presuppose the requirement of per-hop monitoring and propose various monitoring mechanisms to detect the location of the failure [29]. These monitoring mechanisms raise various challenges and may be difficult or costly to implement. The results obtained in this section challenge the assumption that per-hop monitoring is essential in order to induce cooperative behavior or even that it improves the welfare of the participants or the system performance.

However, while monitoring does not affect the expected utility of the participants, several issues arise when comparing the two scenarios ($v1a$ and $v1b$). First, in the absence of monitoring, 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. The tradeoffs between different solution concepts in the context of online environments are discussed in [12, 13].

Second, the payment structure is different. While in $v1a$ the payment to each node depends on its location (i), monitoring provides us with n identical contracts. This difference becomes important if we relax the assumption that the nodes' location is common knowledge. In $v1a$, where nodes in different locations face different contracts, nodes may strategically misreport their location in order to maximize their expected payoff. In contrast, when monitoring is used, the sender should not be concerned about this type of strategic behavior since all n contracts are identical.

One particular mechanism that provides the sender with per-hop information is the *acknowledgment* (ack) mechanism, in which every node sends an ack to the server if it receives the packet. Following the assumption about costless monitoring, we assume that ack transmission is costless and experience no loss. We claimed above that the use of monitoring provides us with dominant strategy equilibrium. However, in this particular monitoring mechanism, forwarding a packet weakly dominates dropping only if node $i+1$ sends an ack if it receives it. Therefore, while the choice to drop or forward is independent of the other agents' *forwarding* behavior, it still depends on their willingness to adhere to the ack protocol. For this reason, the payment scheme needs to address the strategic consideration of ack transmission as well. But if the IR constraint is satisfied and ack transmission incurs no cost, nodes have no incentive to decline ack transmission, since a node is paid only if it sends an ack [4]. Thus, the strategy profile in which all nodes forward and abide by the ack protocol is a Nash equilibrium. However, the monitoring mechanism is vulnerable to collusion among nodes. For example, several nodes may collude to report they have received the packet, thereby gaining the benefit of the high payment without incurring the forwarding cost.

3.3 Priority versus Best-Effort Forwarding ($v2$)

The next scenario we consider is one in which nodes choose between best-effort forwarding ($a = 0$) and priority forwarding ($a = 1$). In contrast to $v1$, a packet may reach the next hop even if the low-effort action is taken. However, if the high-effort action is taken, the loss rate decreases. Consequently, the probability that a packet is successfully transmitted from node i to node $i+1$ is:

$$Pr(x^H_{i\to i+1}|a_i) = 1 - (k - qa_i) \tag{4}$$

$$\text{where: } q \in (0,1] \text{ and } k \in [q,1]$$

We distinguish between the loss rate that is under the control of the node (e.g., buffer overflow), and the loss rate that is beyond the control of the node (e.g., data corruption, network congestion). The former is captured by q, and their aggregated effect is captured by k.

Proposition 3.3 *In $v2$, like in $v1$, the sender derives the same expected utility whether it obtains per-hop monitoring information or not.*

Proof See Appendix B.

While the main result remains the same, the payment structures in $v1$ and $v2$ are different. In particular, while in $v1$ $s^L_i = 0$ at the best contract, in $v2$ $s^L_i < 0$ (see equation 5 in Appendix B). Thus, under limited liability, which requires that $s_i \geq 0$ (essentially, prohibits transfer of money from an intermediate node to the sender), the expected payment to each node will exceed its expected transit cost.

There is an additional difference between $v1$ and $v2$. In $v1$, a sender either signs a contract with all n nodes along the path or with none. This is because one node that drops the packet determines a failure, even if all other nodes potentially forward. In contrast, in $v2$, the payments are used to increase the probability of delivery success, but a success may occur under the low effort actions as well. Therefore, in $v2$ the sender may be better-off signing a contract only with 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. On the other hand, the individual expected payment decreases in i. Therefore, a rational sender who wishes to sign a contract with only m out of the n nodes should sign contracts with the nodes that are closer to the destination; namely, nodes $(n-m+1, ..., n-1, n)$.

3.4 Multi-Path Routing ($v3$)

In this section we consider network scenarios where multiple disjoint paths exist between the sender and the receiver, and the sender can elect to send multiple copies of the packets to maximize the likelihood of delivery. For simplicity, we analyze the network scenario shown in figure 2, but it can be easily generalized to any finite number of parallel disjoint paths. The sender can contract with both A and B, only with one of them, or with none, where the decision depends on the system's parameter, k, c, and $b(x)$. In this section, we consider a sender who chooses to contract with both A and B.

---

[4] The additional condition is: $(Pr(x^H_{i\to i+1}|a_i = 1)s^1_i + Pr(x^L_{i\to i+1}|a_i = 1)s^0_i - c) \geq 0$, which is satisfied.

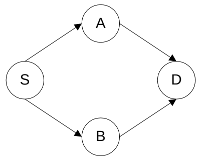

Figure 2: Two parallel disjoint paths from source to destination.

We consider the following two scenarios:

($v3a$) No per-path information: the sender observes whether at least one copy of the packet has reached the destination, but cannot verify the paths through which it arrived. The payments are contingent solely on whether at least one copy arrived.

($v3b$) Per-path information: the sender obtains path-specific monitoring information, and the payments to node i are contingent on successful delivery through the path containing i.

Proposition 3.4 *If the sender elects to send multiple copies of the packets to maximize the likelihood of delivery, it will incur the same expected costs whether it obtains path-specific monitoring information or not.*

Proof See Appendix C.

4 Literature Review

Wireless ad hoc routing has attracted a lot of attention in the context of non-cooperative behavior, including both malicious and selfish behavior. The detection and report mechanisms proposed in [19] route around malicious nodes. While mitigating malicious behavior, they are not effective against selfish behavior, which is the focus of this paper.

In order to mitigate selfish behavior, some approaches require reputation exchange between nodes [4, 5, 6], or simply first-hand observations [3], which are used to reward the well-behaving nodes and punish the misbehaving ones. Other approaches propose payment schemes [7, 18, 21, 29] to encourage cooperation.

The problem of information asymmetry has been studied in the context of BGP routing, using the methodology of distributed algorithmic mechanism design (DAMD) [10, 11]. The focus in DAMD is the design of payment schemes that induce nodes to truthfully reveal their transit costs, thereby enabling nodes to find the optimal paths. However, they do not specify how to enforce the desired forwarding behavior. This problem is considered in [26], which attempts to respond to the challenge of DAMD raised in [12]: “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?” The authors propose various monitoring mechanisms in which nodes audit their neighbours’ behavior. Monitoring is also proposed in [29], where the intermediate nodes send receipts to a third-party entity to indicate that they received the packet. As another example, [21] proposes an audit mechanism in the context of priority forwarding.

The problem of information asymmetry and hidden-action (also known as *moral hazard*) is well studied in the economics literature [8, 15, 20, 22, 28]. [15] 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. [28] shows that this task is made possible when production takes place sequentially, and [2] distinguishes between strategic substitutes and complements, and shows that a principal is better-off under sequential production of strategic complements, but prefers simultaneous production of strategic substitutes.

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 [16, 17, 25].

5 Conclusions and Future Directions

In this work, we identify the problem of hidden-action in multi-hop routing, and show that it is possible to induce cooperative behavior under hidden-action by appropriate contract design. We conclude that monitoring per-hop outcomes in single-path scenarios or obtaining path-specific monitoring information in network scenarios with multiple disjoint paths may not improve the utility of the participants or the network performance.

These results can help influence how system design should proceed with respect to issues such as the desired level of accountability or the need for infrastructure that supports lightweight contracts. In particular, our results show that under some assumptions, it is possible to achieve first-best utility by observing the final outcome solely. In these cases, monitoring may be worthless. However, we may find that under a different set of assumptions, per-hop information may yield some benefit, in which case monitoring mechanisms should be considered.

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

- Uniqueness of the equilibrium. While the proposed payment scheme induces a Nash equilibrium in which all nodes cooperate, it is not the unique equilibrium. We plan to study under what circumstances and mechanisms cooperative behavior can emerge as the unique equilibrium, and highlight the difference between the different scenarios with respect to this issue.
- Sequential contracts. Consider additional models of communication. For example, instead of having the principal communicate with each individual node, nodes can sequentially contract with their successor nodes.
- Risk attitudes. Relax the assumption about the risk neutrality of the agents. While under the standard principal-agent model, a risk-averse agent prevents first-best contracts, it is shown in [22] that risk averse agents can in fact permit

first-best solutions when dealing with teams of multiple agents.

- Correlation of events. We assume that all events are i.i.d. However, the events may be correlated due to different network factors. For example, network congestion or weather conditions may have effect on multiple links concurrently.
- Monitoring. When studying the effect of monitoring under various scenarios, it will become important to consider more realistic monitoring mechanisms and understand the tradeoffs involved.
- Collusion. Monitoring mechanisms in which nodes audit and report the behavior of their peers are vulnerable to various types of collusion. We plan to study what forms of collusion may arise and how they can be mitigated.
- Imperfect information with respect to choice and observability. In this context we plan to study the results of additional solution concepts, such as trembling hand perfection, which suggest that players act given some perturbed beliefs that allow for small mistakes.

6 Acknowledgments

We thank Hal Varian, Scott Shenker, Ion Stoica and Christos Papadimitriou for fruitful discussions. This work is supported in part by the National Science Foundation under grant numbers ANI-0085879 and ANI-0331659.

References

[1] ANDERSEN, D. G., BALAKRISHNAN, H., KAASHOEK, M. F., AND MORRIS, R. Resilient Overlay Networks. In *18th ACM SOSP* (2001).

[2] BANERJEE, A., AND BEGGS, A. Simultaneous versus Sequential Move Structures in Principal-Agent Models. *University of Oxford, working paper* (1997).

[3] BANSAL, S., AND BAKER, M. Observation-based cooperation enforcement in ad hoc networks. *Technical report, Stanford university* (2003).

[4] 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).

[5] 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).

[6] 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).

[7] BUTTYAN, L., AND HUBAUX, J. Stimulating Cooperation in Self-Organizing Mobile Ad-Hoc Networks. *ACM/Kluwer Journal on Mobile Networks and Applications (MONET)* (2003).

[8] CAILLAUD, B., AND HERMALIN, B. *Hidden Action and Incentives*. Teaching Notes. U.C. Berkeley.

[9] CHRISTIN, N., AND CHUANG, J. On the Cost of Participating in a Peer-to-Peer Network. In *3rd International Workshop on Peer-to-Peer Systems* (2004).

[10] 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).

[11] FEIGENBAUM, J., SAMI, R., AND SHENKER, S. Mechanism Design for Policy Routing. In *Yale University, Technical Report* (2003).

[12] 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).

[13] FRIEDMAN, E., AND SHENKER, S. Learning and implementation on the internet. In *Manuscript. New Brunswick: Rutgers University, Department of Economics* (1997).

[14] GOLDSCHLAG, D., REED, M., AND SYVERSON, P. Onion Routing for Anonymous and Private Internet Connections. In *Communications of the ACM* (1999), vol. 42.

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

[16] 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.

[17] 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).

[18] JAKOBSSON, M., HUBAUX, J.-P., AND BUTTYAN, L. A Micro-Payment Scheme Encouraging Collaboration in Multi-Hop Cellular Networks. In *Financial Cryptography* (2003).

[19] 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.

[20] MASS-COLELL, A., WHINSTON, M., AND GREEN, J. *Microeconomic Theory*. Oxford University Press, 1995.

[21] 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).

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

[23] RATNASAMY, S., FRANCIS, P., HANDLEY, M., KARP, R., AND SHENKER, S. A Scalable Content-Addressable Network. In *Proceedings of ACM SIGCOM* (2001).

[24] ROWSTRON, A., AND DRUSCHEL, P. Pastry: Scalable, Distributed Object Location and Routing for Large-Scale Peer-to-Peer Systems. In *Proceedings of IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)* (2001).

[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] 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* (2001).

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

[29] 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).

Appendices

A Derivation of b_t^H

The sender will enter into a contract only if the expected benefit it derives from a successful outcome exceeds its expected implementation cost.

The expected benefit of the sender is:

$$E[b] = b^H Pr(x^H)$$

Consequently, the sender is willing to enter into the contract only if:

$$E[S] \leq b^H Pr(x^H)$$

The expected implementation cost of the sender (total payments) is:

$$E[S] = (1-k)^{n+1} \sum_{i=1}^{n} s_i^H + (1-(1-k)^{n+1}) \sum_{i=1}^{n} s_i^L$$

which, based on equations 2 and 3, yields:

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

Therefore, the threshold benefit, b_t^H, above which the sender will engage in the contract is:

$$b_t^H = \frac{E[S]}{Pr(x^H)} = c\frac{(1-k)^{-n}-1}{k}$$

This expression increases exponentially in the number of nodes, n, and increases polynomially with the loss rate, k.

B Proof of Proposition 3.3

Proof In what follows, we specify the IC and IR constraints when nodes choose between best effort and priority forwarding and no monitoring is used. The constraints are the same as specified in the proof of proposition 3.1, but their values change, based on equation4, to reflect the new version:

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

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

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

$$s_i^L = \frac{(1-k+q)^i c(k-1)}{q} \qquad (5)$$

$$s_i^H = \frac{(1-k+q)^i c(k-1+(1-k+q)^{-n})}{q} \qquad (6)$$

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:

$(IC): (1-k+q)s_i^1 + (k-q)s_i^0 - c \geq (1-k)s_i^1 + (k)s_i^0$

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

and we get the solution:

$$s^0 = \frac{c(k-1)}{q}$$

$$s^1 = \frac{ck}{q}$$

The expected payment of the sender with or without forwarding is the same, and equals:

$$E[S] = \frac{c(1-k+q)(1-(1-k+q)^n)}{k-q} \qquad \blacksquare \qquad (7)$$

C Proof of Proposition 3.4

Per-path information ($v3b$). We denote successful and unsuccessful transmissions from the source to the destination through node i by x_i^H and x_i^L, respectively (e.g., x_B^H means that a packet reached the destination through node B). The payment schedule to node i specifies the payment s_i^H, if a packet reaches the destination through node i, and s_i^L otherwise. We show here the constraints for node A (and the same applies for node B):

$(IC): Pr(x_{A\to D}^H | a_A = 1)s_A^H + Pr(x_{A\to D}^L | a_A = 1)s_A^L - c \geq$
$Pr(x_{A\to D}^H | a_A = 0)s_A^H + Pr(x_{A\to D}^L | a_A = 0)s_A^L$

$(IR): Pr(x_{S\to A}^H | a_S = 1)(Pr(x_{A\to D}^H | a_A = 1)s_A^H +$
$Pr(x_{A\to D}^L | a_A = 1)s_A^L - c) + Pr(x_{S\to A}^L | a_S = 1)s_A^L \geq 0$

Both constraints bind at the optimal contract, and we get the following solution:

$$s_A^L = 0$$

$$s_a^H = \frac{c}{1-k}$$

The expected payment of the sender is:

$E[S] = Pr(x_A^H)Pr(x_B^H)2s^H + Pr(x_A^H)Pr(x_B^L)(s^H + s^L)+$
$Pr(x_A^L)Pr(x_B^H)(s^L + s^H) + Pr(x_A^L)Pr(x_B^L)2s^L$
$= 2c(1-k)$

Binary information ($v3a$). Given the restricted information provided to the sender in this scenario, the payment schedule to node i specifies the payment s_i^H, if at least one packet reaches the destination, and s_i^L otherwise. As before, we show the constraints for node A, and the same applies to node B because of the symmetry of the scenario. All the probabilities presented below are conditional on $a_i = 1 \; \forall i \neq A$, but we omit these parts for presentation clarity.

$(IC): Pr(x_B^H)(s_A^H - c) + Pr(x_B^L)(Pr(x_{A\to D}^H | a_A = 1)s_A^H +$
$Pr(x_{A\to D}^L | a_A = 1)s_A^L - c) \geq$
$Pr(x_B^H)(s^H - c) + Pr(x_B^L)(Pr(x_{A\to D}^H | a_A = 0)s_A^H +$
$Pr(x_{A\to D}^L | a_A = 0)s_A^L - c)$

$(IR) : Pr(x^H_{S\to A})(Pr(x^H_B)(s^H_A - c) + Pr(x^L_B)$
$(Pr(x^H_{A\to D}|a_A = 1)s^H_A + Pr(x^L_{A\to D}|a_A = 1)s^L_A - c)+$
$Pr(x^H_{S\to A})(Pr(x^H_B)s^H_A + Pr(x^L_B)s^L_A)) \geq 0$

We obtain the following solution:

$$s^L = \frac{c(1-k)}{k(k-2)}$$

$$s^H = \frac{c}{1-k}$$

The expected payment of the sender is:

$$E[S] = Pr(x^H)2s^H + Pr(x^L)2s^L = 2c(1-k)$$

which is the same as in $v3b$. ∎