

# Priority Forwarding in Ad Hoc Networks with Self-Interested Parties

Barath Raghavan and Alex C. Snoeren  
University of California, San Diego  
{barath,snoeren}@cs.ucsd.edu

## Abstract

This paper examines the problem of incentivizing autonomous, self-interested nodes in an ad hoc network to cooperate in forwarding packets between non-adjacent nodes. We argue that assumptions about the homogeneity of nodes or their preferences are antithetical to the design philosophy of ad hoc networks. Instead, we postulate that a layered design consisting of policed but unpriced best-effort forwarding and priced priority forwarding effectively separates the issues of policing “misbehavior” from penalizing “greed.” Leveraging existing schemes to provide best-effort service in the face of malicious nodes, we present an initial design of an incentive-based scheme that provides priority forwarding for “well-behaved” nodes.

## 1 Introduction

A key challenge to deploying a practical ad hoc network is ensuring that nodes do, in fact, cooperate to forward packets. Lack of cooperation can come in two flavors: “misbehavior,” where a node does not adhere to the specifications of the protocol, and “greed,” where a node operates in a manner that optimizes a particular local utility function possibly at the expense of other nodes, and thus behaves in a self-interested manner. We argue that while not necessarily distinct, neither case subsumes the other in general, so an efficient ad hoc network should be explicitly prepared to handle both cases. Further, significant benefits can be obtained by decoupling the mechanisms used to prevent each.

While misbehavior is easy to define (although often difficult to detect), greed is a relative concept: one node’s greed is another’s charity. Unfortunately, most existing schemes that attempt to give incentives to self-interested nodes to encourage cooperation assume that all nodes use some fixed utility metric. This fixed utility metric is often based upon a property of the network such as packet loss rate [3].

We observe that different nodes may have differing tolerances for any particular metric. For example, nodes utilizing forward-error correction are likely to be more tolerant of high loss rates than those using an ARQ-based protocol. The existence of multiple axes of greed may invalidate any theoretical performance guarantees provided by an incentive-based system. In particular, the assumption that a single utility metric is sufficient (or even somehow inherent) may lead to the classification of alternatively motivated nodes as malicious, likely violating an oft-held assumption that nodes are merely self-interested and not openly destructive.

Another important characteristic of ad hoc networks is their heterogeneity. Some nodes may be willing to forward traffic, adhere to a pricing system, and maintain accounting state, while others may be either unwilling or incapable of doing such things. For example, consider a node on the edge of a network: it is highly unlikely to be in the position to forward packets for other nodes in the network as it will not serve as anyone’s next hop. Thus, we argue that it is unreasonable to expect that all nodes in an ad hoc network will adhere to or be able to fully participate in a specific pricing protocol. As a result, any incentive-based scheme should not require global participation.

Given the difficulties in defining an incentive-based scheme that copes well with the realities of ad hoc networks, we advocate the use of layering. In particular, we argue that priority forwarding provides an effective mechanism to incentivize self-interested parties independent of their particular metric of greed. Best-effort traffic, on the other hand, need only be policed against misbehavior. Unless a significant fraction of nodes provide best-effort forwarding, the network is largely unusable for all nodes. Therefore, we postulate that most self-interested parties will provide a reasonable degree of best-effort packet forwarding even without incentive-based schemes.

Since the throughput of an ad hoc network is modest at best [8], nodes that behave particularly well can be rewarded with the ability to send priority traffic. Priority traffic is forwarded ahead of any best-effort traffic. We believe that separating the policing of best-effort forwarding from incentive-based priority forwarding provides three distinct advantages compared to existing, monolithic approaches:

1. Nodes not well positioned to earn goodwill for forwarding others' packets are not completely deprived of service.
2. Incentive-based priority forwarding can effectively moderate the behavior of self-interested nodes even in the presence of disinterested or apparently malicious nodes.
3. The existence of a policed best-effort service may obviate out-of-band communication channels to implement virtual currency, enabling the deployment of proposed incentive-based forwarding schemes [13].

We provide initial evidence supporting these claims by describing the implementation of a simple incentive-based priority forwarding scheme assuming the existence of a policed best-effort service and showing that it effectively rewards well-behaved nodes.

## 2 Related work

We divide prior work into two categories: that which attempts to contain the damage caused by malicious nodes, regardless of their reasons for misbehavior, and that which assumes some notion of greed in an attempt to moderate node behavior. We first discuss several instances of each and then show how our layering approach can improve upon an existing technique.

### 2.1 Controlling the misbehaved

Marti *et al.* address the problem of node forwarding misbehavior by introducing a “watchdog” that observes when nodes are misbehaving and a “pathrater” that assigns ratings for paths based on observed node behavior [9]. In their watchdog they take advantage of the broadcast nature of 802.11 to observe the forwarding behavior of neighboring nodes. The impact of such misbehaving nodes is minimized by routing around them; however, a bad node's packets are still forwarded without complaint.

This shortcoming is addressed by the CONFIDANT system [2], which relies upon a concept of node reputation that is earned over time. When suspicious events are observed, nodes send alarms to their “friend” nodes, which take note of the misbehaving node. Misbehaving nodes are denied forwarding privileges and avoided by others through the use of a global black list.

Neither of the two schemes discuss how routing information is propagated or protected. Hu *et al.* describe the Ariadne system [4] which uses cryptographic techniques to prevent subversion of the routing system in use due to false route updates or route request floods.

### 2.2 Satiating the greedy

Alternatively, several authors have studied ad hoc networks in which nodes are modeled as self-interested (but not malicious) entities. One approach, proposed by Buttyan and Hubaux, uses a virtual currency, “nuglets;” each packet is pre-loaded at the origin with nuglets, which are taken as payment by intermediate forwarding nodes (in their Packet Purse model). Their system relies upon tamper-resistant security modules at each node which set up security associations when nodes become neighbors. These modules ensure that nodes cannot overspend.

Zhong *et al.* describe the Sprite system [13] which uses the idea of credit to solve the problem of routing in ad hoc networks of self-interested nodes. The credit system presented therein subsumes all packet routing—the underlying ad hoc routing protocol only exists for packet delivery, not for routing decision making. To handle payment, the system relies upon a centralized Credit Clearance Service (CCS) which handles receipt processing after nodes receive payment from others. They model the receipt collection process as a game and are able to price forwarding to ensure that truth-telling is an optimal strategy for all involved nodes. In their evaluation, they model nodes as power and credit conservative; nodes require sufficient credit and power to send packets. The authors note that Sprite could make use of systems like Ariadne to protect against various other forms of attack; similarly, it is possible that Sprite could be used in a layered architecture such as the one we propose here.

### 2.3 Combining the two

As described, Sprite nodes communicate with the CCS using some out-of-band mechanism to receive payment for receipts they have collected. We contend that this

approach to communication with the CCS is highly impractical. Our layered model, on the other hand, admits a straightforward implementation. Suppose the CCS existed as a node in the ad hoc network. In the original model, the delivery of receipts is not guaranteed; nodes could purposefully drop receipt packets destined for the CCS.

We can ensure receipt delivery in our model: the system has some guarantee of best-effort forwarding, and thus the receipts can *eventually* be delivered. Since techniques exist to prevent double-spending of virtual currencies (*c.f.* Section 3.2), an ARQ-based protocol could be applied. If simple delivery guarantees prove insufficient (perhaps due to greedy behavior of on-path nodes focused on the mistreatment of receipt-carrying packets), a steganographic scheme could be employed to hide the presence of a receipt within other data packets. Our scheme still does not, and cannot, directly address Sprite’s requirement for a central authority, however. In particular, the budget imbalance caused by Sprite’s pricing model seems fundamental to its design.

### 3 Design requirements

Before describing our priority forwarding scheme, we first explicitly state our assumptions and discuss their ramifications.

#### 3.1 Policed best-effort

As part of our layered architecture, we assume some degree of best-effort packet forwarding for the purposes of higher level incentivization for priority forwarding. Instead of designing such forwarding from scratch, we advocate the reuse of features provided by Kyasanur [7] for the MAC layer and Hu [4] for the routing layer.

Policed best-effort service is useful for several reasons. First, since best-effort forwarding exists, not all nodes must be aware of the pricing system for the network to be useful or effective. Second, even with schemes that assume the existence of a centralized authority, no out-of-band mechanism is needed to communicate with the authority, since best-effort delivery guarantees that eventually all packets get through. Third, users poorly situated in the network topology are not denied service despite their inability to earn credit by forwarding others’ packets. Even without any currency they can still have their packets delivered. By contrast, many systems that use virtual currency to incentivize nodes either have to redistribute currency periodically to under-privileged nodes or allow users to go

bankrupt (meaning they won’t be able to send packets until they forward some).

#### 3.2 Virtual currency

As with most other incentive-based forwarding schemes, our priority forwarding mechanism relies on the existence of a secure virtual currency. Many virtual currency systems have been proposed in the past [1, 10] that ensure absolute security of transactions. Existing systems provide guarantees against double spending, currency loss, and fraudulent currency. The micro-payment scheme proposed by Rivest [12] seems especially attractive in an ad hoc environment due to its reduction of communication overhead: it requires “cashing in” only of winning lottery tickets.

A major drawback to existing protocols is their reliance upon centralized nodes or secure hardware for currency management. Similar to a universal definition of greed, these assumptions seem contrary to the spirit of ad hoc networks which inherently have no reliable central authority, nor can they guarantee that nodes carry secure hardware. The alternative is equally daunting, however: how can nodes in an ad hoc network track the currency of others? It is not immediately apparent how to design such a system in the context of a purely distributed system of mutually distrustful, greedy nodes. We leave the design of such a system to future work, but offer a few initial observations.

We postulate that since the currency does not correspond to a precise, discrete commodity, it may suffice for the virtual currency system to make only probabilistic guarantees about both its security and its fairness. In addition, we observe that in a system of greedy nodes, one way to prevent attacks is to make an attack feasible and well known, yet sufficiently expensive; if a node is charged for the traffic generated to mount an attack, and the attack’s potential gain is less than its expense, a self-interested node will not attempt it.

#### 4 A priority forwarding scheme

We now leverage our layered architecture to design an initial system for incentive-based priority forwarding with self-interested parties. Our aim is to ensure nodes that forward priority packets get reasonably compensated for their forwarding, nodes that do not forward packets in a priority fashion are unaffected, and nodes with equal currency and similar topological location receive similar improvements in delivery ratio (or any other metric of interest) for their expenditures.

Our protocol prices priority forwarding such that each node pays a certain price per packet based on the traffic along the forwarding path. We do not intend to provide any strict notion of performance; indeed, we have argued such an effort would likely be futile, as nodes in the network may not participate in the pricing scheme, may “misbehave,” or may have a different notion of self-interested behavior than intended. In order to simplify implementation, we divide time into epochs. Prices change only at epoch boundaries; we have not yet explored the subtle effects of epoch changes upon packets that are in flight.

#### 4.1 Pricing

Define a matrix  $P$  given by the routing algorithm that defines all paths in the network from node  $i$  to node  $j$  by  $P_{ij}$ . Any particular path  $p \in P_{ij}$  consists of a possible forwarding path between  $i$  and  $j$ , exclusive of  $i$  and  $j$  themselves. We define the intrinsic cost of priority forwarding at a node  $k$  as  $c_k$ . We allow  $c_k$  to be set arbitrarily by each node  $k$ , except to require that there is no cost at nodes that do not support priority forwarding ( $c_k = 0$ ).

Define  $T_k$  to be the number of packets received at  $k$  in the previous epoch<sup>1</sup>. Each node receives payment,  $m_k$ , for forwarding a priority packet in proportion to the total amount of traffic (both priority and best effort) forwarded by  $k$  in the previous epoch,

$$m_k = \beta \cdot T_k, \quad (1)$$

where  $\beta$  is chosen as described below.

The per-packet cost to send a priority packet from  $i$  to  $j$  along a given path  $p \in P_{ij}$  is the sum over all node costs along the path, denoted by  $c(p)$ :

$$c(p) = \sum_{k \in p} m_k \quad (2)$$

The total amount due at node  $k$  in a particular epoch is then simply  $m_k$  times the number of priority packets forwarded by  $k$  and denoted  $M_k$ . A node then has a utility function for forwarding a packet,  $u_k$ :

$$u_k = m_k - c_k. \quad (3)$$

We wish to ensure that  $u_k \geq 0$ , and thus set  $\beta \geq c_k/T_k$  for all  $k$ .

<sup>1</sup>We assume an epoch to be significantly longer than the one-way delay in the network.

#### 4.2 Forwarding

For each packet received at a node  $k$ , the node has precisely two options: it can drop it, or it can forward the packet to node  $i$ , the next hop along the packet’s path  $p$ , in a priority fashion. For now, we will assume a node makes the same decision for all priority packets in a given epoch. For each priority packet that it forwards,  $k$  must include payment equal to the currency previously attached to the packet less  $m_k$ , thus taking a payment of  $m_k$ . In order to earn this payment,  $k$  must send the priority packet before any best-effort packets in its queue. Proper forwarding behavior (the packet is indeed forwarded before any best-effort traffic, is still marked as best-effort, and includes the appropriate amount of currency) is enforced by node  $i$  by promiscuously observing  $k$ ’s transmissions. If  $k$  neglects its forwarding duties, its payment is voided.

While it is possible for prices to become inflated due to high intrinsic costs at certain nodes, the prices will become proportionally high for all nodes. To bootstrap the protocol, we assume each node  $k$  is given some initial amount of currency  $\lambda_k$ . Each node is now faced with the problem of price discovery: how do nodes determine the prices to forward through other nodes? We suggest that since protocols such as DSR [5] and AODV [11] require a route request, price discovery for a route can easily be piggybacked with a route request. Since the price per packet at a node is proportional to the number of packets seen at that node in the last epoch, our pricing scheme provided with an appropriate objective function should fall within standard pricing stability requirements [6]. Ensuring accurate dissemination of pricing information is an area for future work.

#### 4.3 Simulation

To evaluate our proposed priority scheme, we study several scenarios using the network simulator ns. In particular, we examine the behavior of prices and currency at nodes with priority pricing, the gains of a priced priority forwarding system, and the differentiated gains of nodes with different initial budgets.

All simulations use the fixed topology shown in Figure 1. We construct the topology shown due to its symmetry (for comparison of “similar” nodes) yet non-trivial forwarding paths. Routing is conducted using the AODV protocol; routing messages are forwarded as priority traffic but are ignored by the pricing system. Traffic is sent from eight source nodes to eight sink nodes

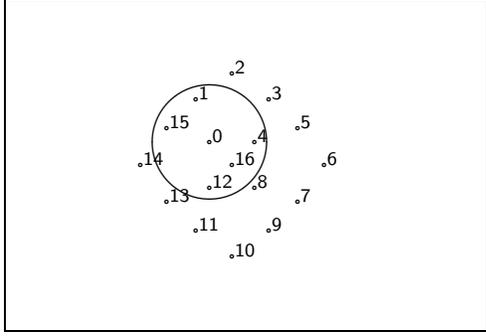


Figure 1: Topology with the radio range for Node 0 shown.

(nodes 0–7 send to 8–15, with node 16 idle but capable of forwarding). In each simulation, we vary the load on the network from two packets/sec/source (8 Kb/sec aggregate) to 20 packets/sec/source (80 Kb/sec aggregate). Each simulation simulates 200 seconds of packet transmission followed by 100 seconds of cool down. Node prices are recalculated each second, with initial prices set to zero. All graphs represent the average of five simulation runs. In an attempt to capture the most general form of greed, we make the assumption that a packet will be sent as priority whenever a node has money to do so, obviating the need to consider any specific settings for  $c_k$  in the simulation.

#### 4.3.1 Pricing fairness

One goal of our system is for currency to have equal value to all nodes. That is, the improvement in delivery ratio obtained by spending any fixed amount of currency,  $\lambda$ , should be the same across all similarly situated nodes. We examine the fairness of our currency by considering a greedy policy whereby a node sends its traffic as priority whenever money is available, and resorts to best-effort if insufficient currency is available. Using the results of previous simulations (not shown), we fix  $\lambda$  lower than the desired currency for every node. In other words, no node has sufficient funds to send all its traffic as priority traffic.

We simulate the behavior of two different nodes that are situated at similar places on a pentagon: nodes 1 and 7. Since the topology is symmetric across several axes, this particular choice of node pair is likely representative of other pairs' behavior. Both 1 and 7 receive the same fixed amount of initial currency. Figure 2 shows the increase in delivery ratio versus an identical simulation (not shown) where both nodes send all traffic as

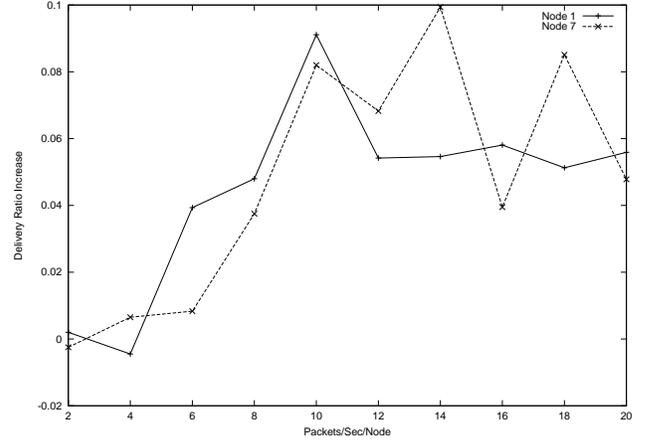


Figure 2: Increase in delivery ratio with fixed currency. The absolute delivery ratio is initially near 1; hence, no improvement is possible at low packet rates.

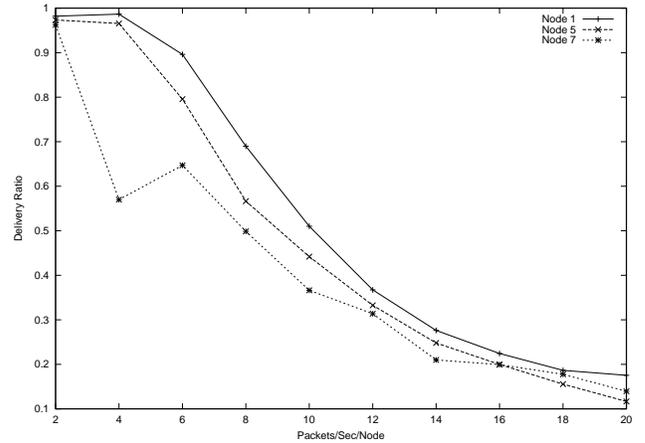


Figure 3: Delivery ratio with varied currency.

best effort. Through the 200 second active simulation period, the nodes turn on and off prioritization several times as they earn money from forwarding packets and spend it sending priority packets.

#### 4.3.2 Marginal utility

Next we show that our scheme is capable of providing different levels of service to nodes with different initial currencies. Figure 3 shows a simulation with nodes 1, 5, and 7 requesting priority delivery; all three nodes are similarly situated in the topology. The three nodes receive roughly linearly decreasing amounts of money, with node 1 receiving the most initial currency and node 7 receiving the least. In this scenario, since the nodes are requesting priority simultaneously, they compete for resources; their utilizations are directly attributable to the amount of initial currency given to each.

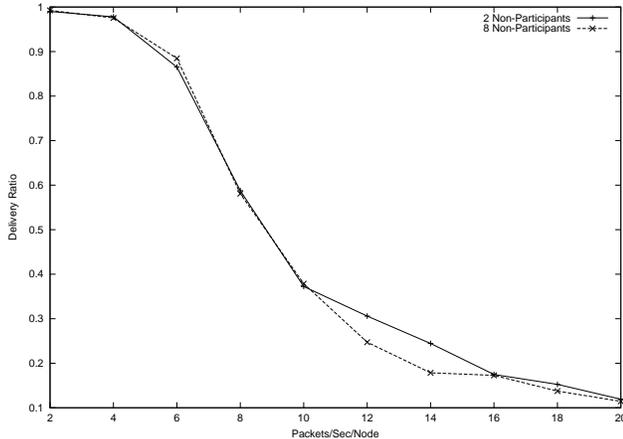


Figure 4: Delivery ratio with non-participating nodes.

### 4.3.3 Partial deployment

The feasibility of partial deployment of a pricing system is one of our primary arguments for layered best-effort and priority priced forwarding. We show in simulation that this can be achieved with our pricing system. Figure 4 shows the delivery ratio with node 2 sending priority traffic with two differing degrees of partial deployment: two or eight centrally located nodes do not participate in the pricing system in the plot shown. The results show that performance gains due to prioritization decrease slightly as more nodes in the system do not cooperate with the pricing system, especially under load, but that the nodes still receive some benefits even with partial deployment. It is also worthwhile to note that this delivery ratio curve for node 2 is similar to those observed previously for nodes 1, 5, and 7.

## 5 Conclusions and future work

We argue that a priced priority forwarding scheme built upon a policed best-effort forwarding system affords more flexibility with respect to a heterogeneous user population while still enabling service differentiation and various degrees of fairness. In both our study of previous work and the design of our proposed system, we identify several problems that must be addressed by any ad hoc network pricing system.

Chief among these, we believe that the design of a virtual currency system that takes into account the innate characteristics of ad hoc networks is an important next step toward practical deployment of incentive-based forwarding protocols. Such a currency system would have to assure no double spending, forgery, or currency loss, yet keep limited state without reliance

upon trusted hardware or centralized authorities. We hypothesize that considering relaxed, statistical guarantees might be a necessary step in developing such a distributed currency system.

We have similarly argued that the consideration of a universal utility metric is misguided—it is likely that certain utility metrics are inherently more useful than others in the context of forwarding incentives. We seek to explore, from both an experimental and theoretical standpoint, the relationships among different greed metrics, the behavior of nodes, and the resulting system's equilibria.

## References

- [1] M. Bellare, J. Garay, C. Jutla, and M. Yung. VarietyCash: a multi-purpose electronic payment system. In *Proceedings of 3rd Usenix Workshop on Electronic Commerce*, pages 9–24, August 1998.
- [2] S. Buchegger and J.-Y. L. Boudec. Performance analysis of the CONFIDANT protocol. In *Proceedings of ACM MobiHoc*, pages 226–236, June 2002.
- [3] L. Buttyan and J.-P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *ACM/Kluwer Mobile Networks and Applications*, 2003.
- [4] Y.-C. Hu, A. Perrig, and D. B. Johnson. Ariadne: a secure on-demand routing protocol for ad hoc networks. In *Proceedings of ACM MOBICOM*, pages 12–23, September 2002.
- [5] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, pages 153–181. Kluwer Academic Publishers, 1996.
- [6] Y. A. Korilis and A. Orda. Incentive compatible pricing strategies for QoS routing. In *Proceedings of IEEE INFOCOM*, pages 891–899, March 1999.
- [7] P. Kyasanur and N. Vaidya. Detection and handling of MAC layer misbehavior in wireless networks. In *Proceedings of the International Conference on Dependable Systems and Networks*, June 2003.
- [8] J. Li, C. Blake, D. S. De Couto, H. I. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *Proceedings of ACM MOBICOM*, pages 61–69, July 2001.
- [9] S. Marti, T. J. Giuli, K. Lai, and M. Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *Proceedings of ACM MOBICOM*, pages 255–265, August 2000.
- [10] G. Medvinsky. *A Framework for Electronic Currency*. PhD thesis, USC, 1997.
- [11] C. E. Perkins and E. M. Royer. Ad hoc on-demand distance vector routing. In *Proceedings of IEEE WMCSA*, pages 90–100, February 1999.
- [12] R. L. Rivest. Electronic lottery tickets as micropayments. In *Financial Cryptography*, pages 307–314. Springer, 1997.
- [13] S. Zhong, J. Chen, and Y. R. Yang. Sprite: A simple, cheat-proof, credit-based system for mobile ad-hoc networks. In *Proceedings of IEEE INFOCOM*, March 2003.