HAL open science

Modelling Incentives for Collaboration in Mobile Ad Hoc Networks

Jon Crowcroft, Richard Gibbens, Frank Kelly, Sven Ostring

To cite this version:

Jon Crowcroft, Richard Gibbens, Frank Kelly, Sven Ostring. Modelling Incentives for Collaboration in Mobile Ad Hoc Networks. WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks, Mar 2003, Sophia Antipolis, France. 10 p. inria-00466747

HAL Id: inria-00466747

https://inria.hal.science/inria-00466747v1

Submitted on 24 Mar 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Modelling Incentives for Collaboration in Mobile Ad Hoc Networks

Jon Crowcroft†, Richard Gibbens†, Frank Kelly‡, Sven Östring†
Computer Laboratory†, *Statistical Laboratory*‡,
University of Cambridge.
{Jon.Crowcroft, Richard.Gibbens, Sven.Ostring}@cl.cam.ac.uk†
F.P.Kelly@statslab.cam.ac.uk‡

Abstract

This paper explores a model for the operation of an ad hoc mobile network. The model incorporates incentives for users to act as transit nodes on multi-hop paths and to be rewarded with their own ability to send traffic. The paper explores consequences of the model by means of fluid-level simulations of a network and illustrates the way in which network resources are allocated to users according to their geographical position.

1 Introduction

There are good reasons why nodes in a mobile ad hoc network, that lacks the networking infrastructure which has been deployed through the investment of a telecommunications corporation, would prefer not to cooperate. When nodes *do* cooperate, they form the necessary ad hoc infrastructure that makes multi-hop communication achievable, allowing traffic from a node to reach destinations that would either require a significant amount of transmission energy using single hop communication, or simply not be possible without routing the traffic through other nodes. However, this means that nodes must be willing to forward traffic for other nodes, and in this way expend energy without receiving any direct gain from doing so. If a node only considers its own short-term utility, then it may not choose to participate within the network.

Thus, the concept of introducing incentives for collaboration into the architecture of this type of network is a important step, and one which allows us to consider the dynamics of the cooperation and preferences of nodes within a system. This leads us naturally to the use of pricing mechanisms, which have found application in rate control in wireline networks [6, 3, 8] and resource control in wireless networks [7, 9, 11]. The difference in this situation is that nodes recover costs, associated with energy losses and traffic loading at a particular node, through the credit arising from pricing mechanisms. This has been shown to stimulate cooperation within ad hoc networks [1]. Determining energy-efficient routes is also an important consideration in ad hoc networks [2, 12], and pricing mechanisms provide the means of guiding a system to its optimal operating point.

In this paper, we specifically consider the issue of how prices can be determined automatically by the ability of nodes to pay the costs for transmitting traffic, and the routes that are subsequently used. We show that cooperation is a natural outcome that emerges from incentives created by the pricing mechanisms. We further study the way that the mobility of the users affects the sharing of resources.

2 System Description

We model our network as a set $\mathcal{N}$ of mobile nodes that are equipped with directional, wireless antennas, with $N = |\mathcal{N}|$ being the number of nodes. Note that in this paper we use the terms "node" and "user" interchangeably to refer to the same entity in the network, where a unique integer index can be used to distinguish between two distinct entities. The difference between the two terms is subtle, where the term "user" is more closely associated with a person who desires to send traffic to other users in the network

and pays congestion costs for doing so. In comparison, the term "node" has topological meaning, in terms of position, velocity, capacity constraints and routing. However, these properties all belong to a single entity, so that the terms are not completely distinct, and the respective term is used in this paper where the natural meaning of the word is appropriate.

Amongst the set $\mathcal{N}$ of nodes, there is a set $\mathcal{S}$ of sources that have destinations $\mathcal{D}$ to send traffic to. To do this, a set of routes between each source and destination pair has been determined, where a route $r \subset \mathcal{N}$ is a subset of the nodes. These routes can be determined using routing protocols like AODV [10] or DSR [5]. Within the set $\mathcal{R}$ of all the possible routes within the network, we can identify $\mathcal{R}^{\mathcal{S}}(s)$ as the subset of routes that originate at source s, and $\mathcal{R}^{\mathcal{D}}(d)$ as the subset of routes that terminate at destination d.

With regards to traffic flow at a specific point in time, each source is transmitting a total amount of traffic x_s, which may be split between the routes $r \in \mathcal{R}^{\mathcal{S}}(s)$. Optimisation of traffic flows, from a single source using multiple routes, has been considered previously [6, 12, 13]. The traffic flow along a particular route r is given by y_r, where $y_r \geq 0$ and

$$x_s = \sum_{r \in \mathcal{R}^{\mathcal{S}}(s)} y_r. \tag{1}$$

We consider that a node is restricted by only having one transceiver. We also observe that a node has limits on its capacity to transmit or receive, where the capacity limit is defined by its spectrum allocation and medium access protocol. The total flow constraint can be modelled by calculating the total capacity usage, where traffic that is forward by a node must be both received and transmitted:

$$c_j = \sum_{\substack{r: j \in r, \text{ and} \\ r \in \mathcal{R}^{\mathcal{S}}(j) \cup \mathcal{R}^{\mathcal{D}}(j)}} y_r + \sum_{\substack{r: j \in r, \text{ and} \\ r \notin \mathcal{R}^{\mathcal{S}}(j) \cup \mathcal{R}^{\mathcal{D}}(j)}} 2y_r, \tag{2}$$

and constraining the capacity usage as follows:

$$c_j \leq C_j, \quad \forall j \in \mathcal{N}, \tag{3}$$

Notice that this constraint does not fully capture the interference issue that arises in wireless networks, but it is a simplification that ensures that a node cannot receive traffic from, or transmit traffic to, two of its neighbours simultaneously.

A key issue within mobile ad hoc networking is energy efficiency, and this can be achieved through traffic management and optimal routing of traffic flows. The energy consumed per unit flow when transmitting traffic from node i to node j is represented by the variable $e_{ij}^{(tx)}$. Receiving traffic also consumes energy, though is independent of the node from which the traffic was transmitted, so we represent this energy consumption by the constant $e^{(rx)}$. Note that variables $e_{ij}^{(tx)}$ can vary with time, representing the mobility of our system. Also, if a node j cannot be reached from node i, then $e_{ij}^{(tx)} = \infty$. Finally, we use the notation that f_{ri} is the node that i will forward traffic to, when using route r.

When considering a specific node j, the power consumed by the node is:

$$\gamma_j = \sum_{r \in \mathcal{R}^{\mathcal{S}}(j)} y_r e_{jf_{rj}}^{(tx)} + \sum_{r \in \mathcal{R}^{\mathcal{D}}(j)} y_r e^{(rx)} + \sum_{\substack{r: j \in r, \text{ and} \\ r \notin \mathcal{R}^{\mathcal{S}}(j) \cup \mathcal{R}^{\mathcal{D}}(j)}} y_r \left(e^{(rx)} + e_{jf_{rj}}^{(tx)} \right). \tag{4}$$

Power consumption is constrained at a node, due to the rate of discharge of the node's battery. This leads us to the following power constraint at each individual node:

$$\gamma_j \leq \Gamma_j, \quad \forall j \in \mathcal{N}, \tag{5}$$

where Γ_j depends on the specification of the node's power supply.

3 Model

3.1 Dual algorithm for flow allocation

Our model for flow allocation builds on that described in [6] for route selection. Suppose that each user s has a parameter $w_s(t)$, known as the willingness-to-pay parameter, and that the user adjusts its total

flow rate on route r (where $r \in s$) as a function of time according to the expression

$$x_s(t) = \sum_{r \in s} y_r(t) = \frac{w_s(t)}{\min_{r \in s} \sum_{j \in r} \mu_{jr}(t)} \tag{6}$$

with $y_r(t)$ only being positive on routes r that attain the minimum in the denominator. This model closely reflects what will occur in a mobile ad hoc network, as the lowest cost paths will be selected in practice. Prices along the routes are defined according to the following equation:

$$\mu_{jr}(t) = \begin{cases} e^{(tx)}_{jf_{rj}} \mu^P_j(t) & + & \mu^B_j(t), & j \text{ is the source node on route } r, \\ \left(e^{(rx)} + e^{(tx)}_{jf_{rj}}\right) \mu^P_j(t) & + & 2\mu^B_j(t), & j \text{ is a transit node for route } r, \\ e^{(rx)} \mu^P_j(t) & + & \mu^B_j(t), & j \text{ is the destination node on route } r. \end{cases} \tag{7}$$

where the congestion prices $\mu^P_j(t)$ and $\mu^B_j(t)$, for power and bandwidth respectively, are dynamically adapted according to the equations:

$$\frac{d}{dt}\mu^B_j(t) = \frac{\kappa \mu^B_j(t)}{C_j} \left(c_j(t) - C_j\right), \tag{8}$$

and

$$\frac{d}{dt}\mu^P_j(t) = \frac{\kappa \mu^P_j(t)}{\Gamma_j} \left(\gamma_j(t) - \Gamma_j\right). \tag{9}$$

The dependence of the right-hand sides of (8) and (9) on both the current price and the respective capacity are an attempt to scale the dynamics of the prices in a network with widely differing prices and capacities. The overall effect of this dual algorithm is, under stable operation, to allocate flows to routes for each user in such a way that the traffic for a given user s, say, faces congestion costs at the rate of $w_s(t)$ per unit time. Global stability of the system (6–9) can be established by the construction of an appropriate Lyapunov function [6], in the case where the network structure is static. We investigate here a model where both the network structure and the set of sources is varying over time.

3.2 Balancing the congestion costs

So far, our model produces a traffic allocation across possible routes determined by the willingness-to-pay parameters, $w_s(t)$, for each user. We now seek to provide an incentive for a user j, say, to act as a transit node for other user's traffic by supposing that user j receives a notional credit for the congestion costs it incurs from each individual source with routes passing through j. The credit thus accumulated can then be re-cycled as payment to other nodes which act as transits for traffic originating at the resource. In this way, users will have the strongest incentive to act as transits where there is the greatest excess demand for traffic, since they earn the most in transit fees. Note that we consider a node to represent a composite resource, having both capacity and energy resources.

We suppose that each user maintains a credit balance, $b_s(t)$, which receives an initial endowment of 1 when the user arrives into the system, where we here identify a node with the user, s, say, whose routes originate at that node. The user's credit balance is then adjusted by transferring credit equal to the congestion costs to each of the downstream resources. Each user will seek to control their credit balance, $b_s(t)$, and we envisage them doing so by dynamically adjusting their willingness-to-pay parameters, $w_s(t)$, according to the level of their credit balance by following a rule of the form: $w_s(t) = \alpha_s b_s(t)$ for some parameter $\alpha_s > 0$. In this way, the user's sending rates would become coupled with their credit balance and they would thereby naturally reduce their sending rate whenever their credit balance was low.

The credit balance itself is discounted, over time,

$$\frac{db_s}{dt} = -\beta \left(b_s(t) - 1\right) - w_s(t) + \sum_{r: s \in r} y_r \mu_{sr}(t) \tag{10}$$

where β is a small positive constant. This will tend to keep the sum of the credit balances $b_s(t)$ over sources $s \in \mathcal{N}$ (the total credit in the system) near the population size N: when a user leaves the system (with their credit lost to the system), the total credit in the system will adjust towards the size of the remaining population.

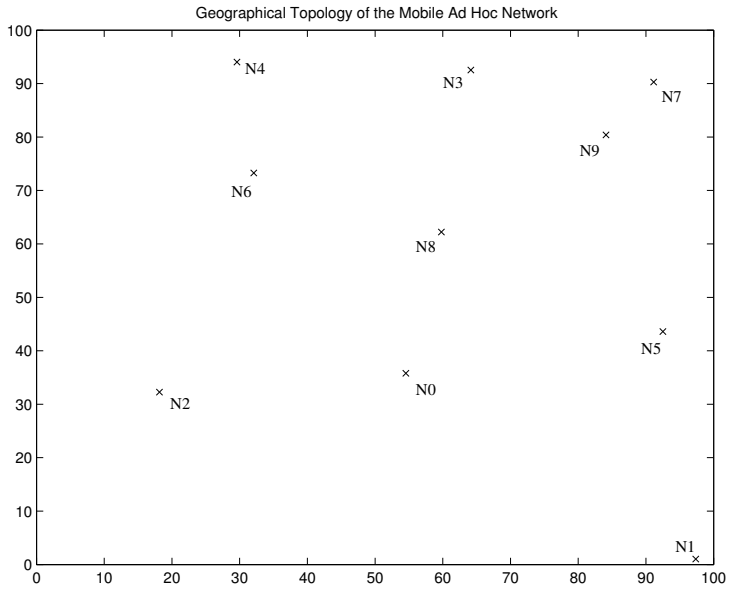

Figure 1: Topology of the mobile ad hoc network.

4 Simulations

The dynamics of the system described in Section 3 are illustrated here using a simulation model. In particular, we demonstrate the stability of prices at nodes and their credit balances. With regards to performance, we also investigate the throughput of the system. Certain dynamics of the system are also studied, including the arrival and departure of users from the system and how this affects the total credit. Finally, we consider how user mobility affects their individual throughput and also how it contributes to the overall system throughput.

The initial network topology includes ten randomly placed nodes in a geographical area of $100m$ by $100m$, as shown in Figure 1. While the topology has been generated by randomly locating nodes using a uniform distribution, it has the features of higher clustering of nodes towards the top-right corner, a node closely situated within the geographical centre of the network (node $N8$), and nodes with higher geographical isolation (particularly nodes $N1$ and $N2$). This allows us to study a set of nodes with diverse geographical locations and topological relationships.

Each node is equipped with a single transceiver with range 56 metres, which defines the neighbours that it has within the network. Notice that for nodes in the centre of the network, this means that nodes have a large set of neighbours and hence have a higher number of routes to choose from, in order to send traffic to a particular destination. Nodes, such as $N1$ and $N2$, have only a few neighbours, and so can only select routes from a smaller set of possible paths.

With regards to traffic model, we assume that a particular user wants to communicate with all other nodes in the network, but establishes a connection with a randomly selected recipient in the network at particular points in time, where connection durations are exponentially distributed. Once the connection terminates, the user remains idle for an exponentially distributed period. When a user initiates a connection with another user, it determines the lowest cost route and then continues to use that particular route for the duration of the connection. Notice that this is a departure from the model described in Section 3, where users continually monitor all available routes to the recipient user and always route traffic through the route with minimum cost. However, this departure from the model is a realistic one, because we want to minimise the amount of routing information that has to be distributed within the network. When using one of the proposed ad hoc routing protocols, such as AODV [10] or DSR [5], it is reasonable to assume that the integrity of routes will need to be checked before routing a stream of packets along a particular path. However, it is unlikely that nodes will continuously monitor all paths at the granularity level of transmitting each packet. The consequence of this departure from the model is that the system will not achieve optimal performance, but there is a trade-off between optimality and the overhead involved in continuously monitoring the prices of other routes to the destination. Another

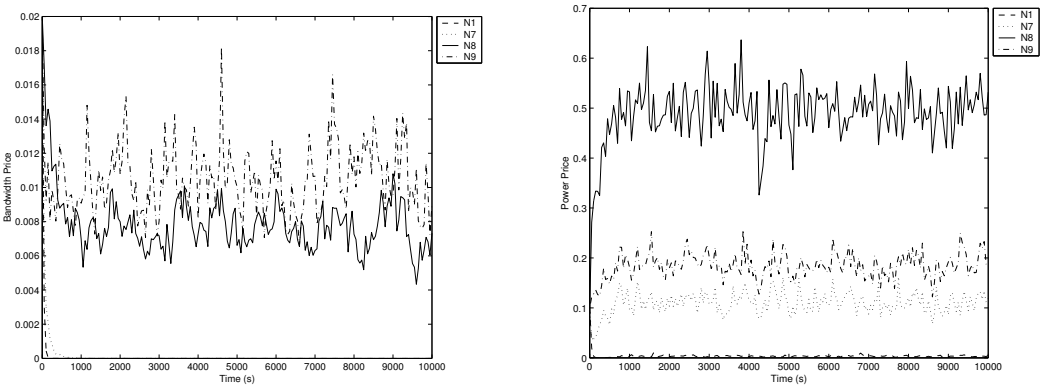

(a) Bandwidth Prices

(b) Power Prices

Figure 2: Bandwidth and power prices curves of four representative nodes ($N1$, $N7$, $N8$ and $N9$).

advantage of this approach is that route-flapping is avoided, which may occur if the price of another route drops below the route currently used, and then a user begins to frequently swap traffic between these two routes.

4.1 Static Network: Stability of Price and Credit Balances

To demonstrate the stability of the system, we simulate a static network topology for $10,000s$ where the mean duration of a connection is $0.5s$, and a user is idle for a mean period of $0.5s$ after completing a connection. Each user updates their prices every $0.01s$. The system parameters used in the system are set as $\alpha_s = 0.3$, $\beta = 0.01$ and $\kappa = 0.05$. The bandwidth capacity is set to $C = 10$ for all nodes in the network, while the maximum power is $\Gamma = 0.5$. The energy parameters associated with transmitting and receiving traffic are given by $e_{ij}^{(tx)} = 10^{-4}\|z_i - z_j\|_2^{1/2}$ and $e^{(rx)} = 10^{-3}$, with z_i and z_j being the geographical position of nodes i and j respectively.

The prices and credit balance of four representative nodes in the network are shown in Figure 2. Node $N1$ has been selected, as it is the most extreme node in the network, while node $N7$ is an extreme node with nodes $N3$ and $N9$ in close proximity. The prices of the node nearest the centre of the network, namely $N8$, have also been plotted, together with those of node $N9$, which is also frequently used as a transit node.

It can be observed from these plots that each price stabilises about a mean value, hence providing evidence that the overall system is stable. It should be noted that this occurs with the sub-optimal routing policy that minimum cost routes are only selected when connections are established. A second observation is that prices for node $N1$, which is on the edge of the network, all decay rapidly to zero. This is because no routes are selected which use $N1$ as a transit node, and the only flows which consume bandwidth or power resources at this node are those originating or terminating at $N1$.

It is also interesting to compare the prices of nodes $N8$ and $N9$. The bandwidth price is the highest for $N9$, while $N8$ has the highest power price. The reason for this is that while node $N8$ is the closest to the centre of the network, distances to its neighbouring nodes are all relatively high. Hence, as more power is consumed by $N8$ in transmitting to other nodes, the power price will be driven up. In comparison, $N9$ is not near the centre of the network, yet is close to nodes $N3$ and $N7$, and will be carrying larger amounts of traffic for these nodes and for other nodes that route traffic around this cluster of nodes. Hence, its capacity usage will be reasonably high, as reflected by its bandwidth price.

The credit balances and throughputs, for the same nodes, are plotted in Figure 3. Throughput is determined by logging the accumulative traffic originating from the node in $50s$ intervals. Once again, note that these quantities stabilise around their individual mean values. These mean value for each node's credit balance is largely dependent on their geographical location within the network. As would be expected, node $N8$ maintains the highest credit balance, as it will be carrying a large amount of transit traffic. In addition, $N8$ will be charging high power prices for doing so, and thus accruing significant

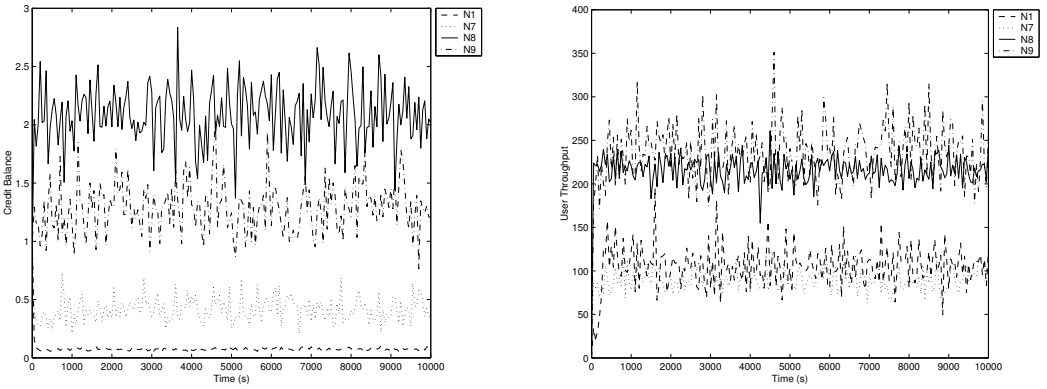

(a) Credit Balances

(b) Throughputs

Figure 3: Credit balances and throughput curves of four representative nodes ($N1$, $N7$, $N8$ and $N9$).

credit in the process. As the location of a node gets closer to the edge of the network, its credit balance is seen to decrease.

4.2 Dynamic Network: Arrivals and Departures of Users

Having demonstrated that the system does stabilise for a static network using simulations, we now investigate a dynamic network where users can join and leave the network, depending on decisions based on conditions that are external to the network. Consequently, we model the arrivals and departures as a random process. In particular, we consider the arrivals as a Poisson process and the "lifetime" of a user in the network is exponentially distributed. The location of user, when it arrives, is randomly distributed within the $100m$ by $100m$ square area in which the network is situated.

We are concerned with the consequence to the total credit in the system when users join or leave the network. A user who arrives will always increase the total credit by one, due to its initial credit endowment. However, the situation is not the same when a user leaves the system, as the user may have accumulated a large amount of credit from other users, because of its ability to act as a transit node. When such a node leaves the network, the total credit will decrease by more than one. Otherwise, if a user has spent most of its credit paying transit fees for traffic routed through other nodes, then the total credit will not decrease significantly. In both cases, at the instant when the user departs, the total credit will not reflect the true number of users in the system.

The parameter β has been introduced into the model in (10) to discount the balance of users, over time, so that the total credit in the system will adjust to the true number of users in the system. This property of the system is shown in Figure 4. The mean arrival rate of users is 3.6 users per hour, and each user remains in the system for a mean period of 16.7 minutes. Figure 4 shows that the total credit in the system tracks the number of users in the system. The rate of decay to the actual number of the system is defined by the value of β.

4.3 Mobile Network: User Prices and Overall Throughput

The final objective of this paper is to study the effect of mobility on the performance of our ad hoc system, where nodes have incentives to collaborate. Returning to the original topology considered earlier in Figure 1, the most extreme node $N1$ is mobilised, and follows the path, shown in Figure 5, through the geographical centroid of the static network consisting of the remaining nodes. We observe the performance of the system, as the $N1$ moves across the network and reaches the other edge of the network by the end of the simulation which is run for $10,000s$. To reach this final location, the velocity of $N1$ is set to $(-0.0074, 0.0126)m/s$.

As it approaches the centre of the network, $N1$ will be used more frequently as a transit node to carry traffic between other nodes, and this can be observed from the increase in both the bandwidth and power

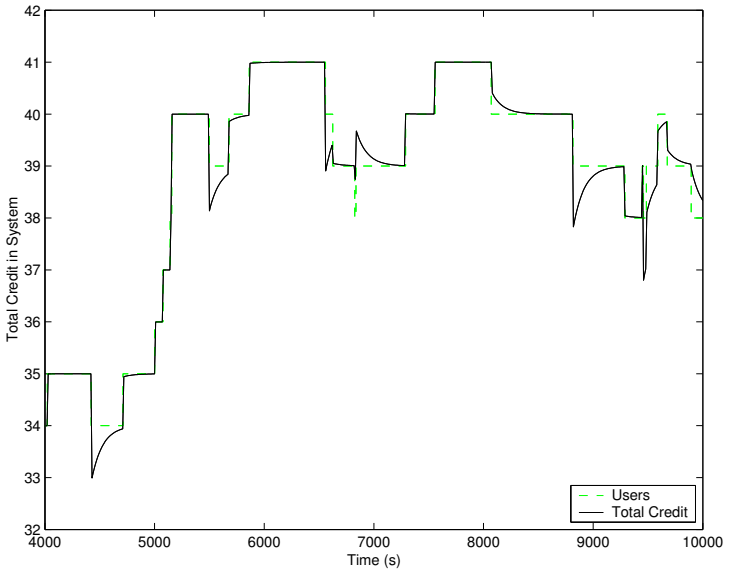

Figure 4: The number of users and the total credit in the dynamic network.

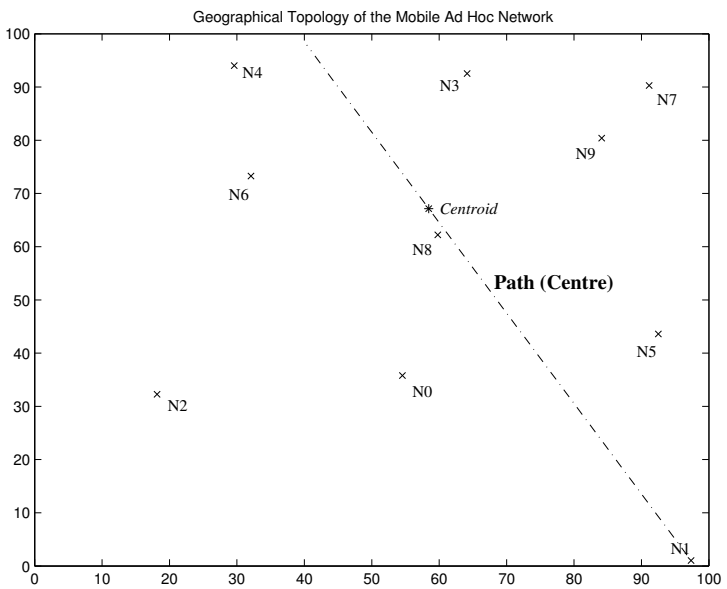

Figure 5: Paths of nodes through the network.

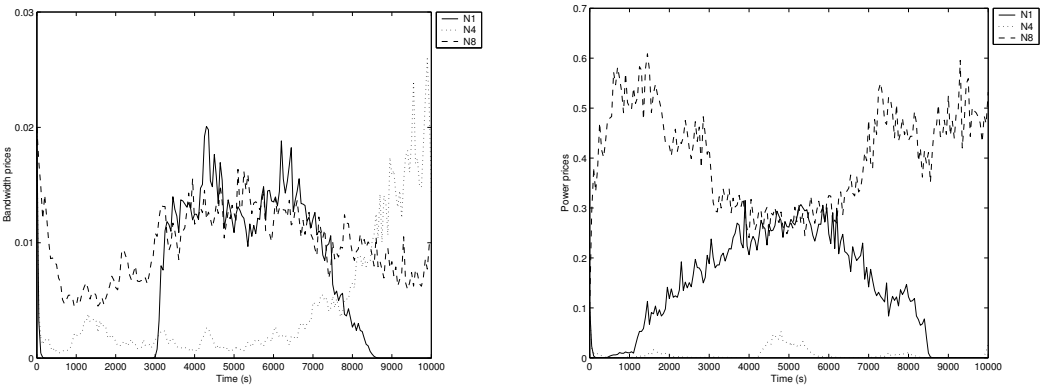

(a) Bandwidth Prices

(b) Power Prices

Figure 6: Bandwidth and power prices curves of the mobile node $N1$, and two stationary nodes $N4$ and $N8$.

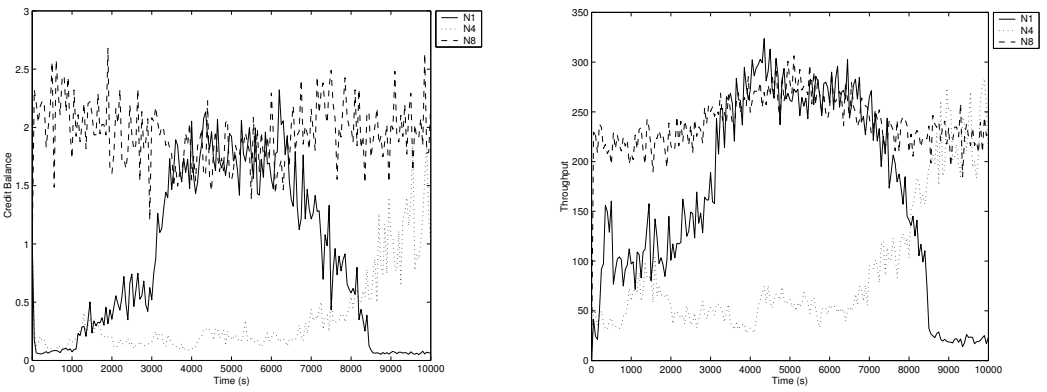

(a) Credit Balances

(b) Throughputs

Figure 7: Credit balances and throughput curves of the mobile node $N1$, and two stationary nodes $N4$ and $N8$.

price of node $N1$ in Figure 6. At the same time, other nodes will now have a choice of sending traffic through either $N8$ or $N1$, when both nodes are near the centre of the network, so the effect of $N1$ moving to the centre is to reduce substantially reduce the power price of $N8$. As node $N1$ moves away from the centre of the network, these effects on the node prices subside.

The increase in prices associated with node $N1$, when it is near the centre of the network, and its increased traffic load which it forwards for other nodes, means that its credit balance also grows, as shown in Figure 7(a). This increases the ability of $N1$ to generate traffic, as its willingness-to-pay is related to its credit balance. Consequently, its throughput increases, as can be observed in Figure 7(b). Due to the competition between $N1$ and $N8$, the credit balance of node $N8$ decreases slightly. However, it should be noted that while $N8$'s credit balance decreases, its overall throughput increases. This is principally due to the fact that $N1$ is now much closer to $N8$, and so the actual cost of sending traffic to $N1$ becomes substantially less. This results in increasing the traffic load between $N1$ and $N8$, and so the bandwidth price of node $N8$ increases accordingly. This increase in throughput and bandwidth price, when a node moves closer to another particular node, is also observed in Figures 6(a) and 7(b) when $N1$ moves away from the centre of the network and closer to node $N4$.

The remaining question is whether the mobility of node $N1$ through the centroid of the network effectively increases the overall throughput of the system. Figure 8 shows that the overall throughput

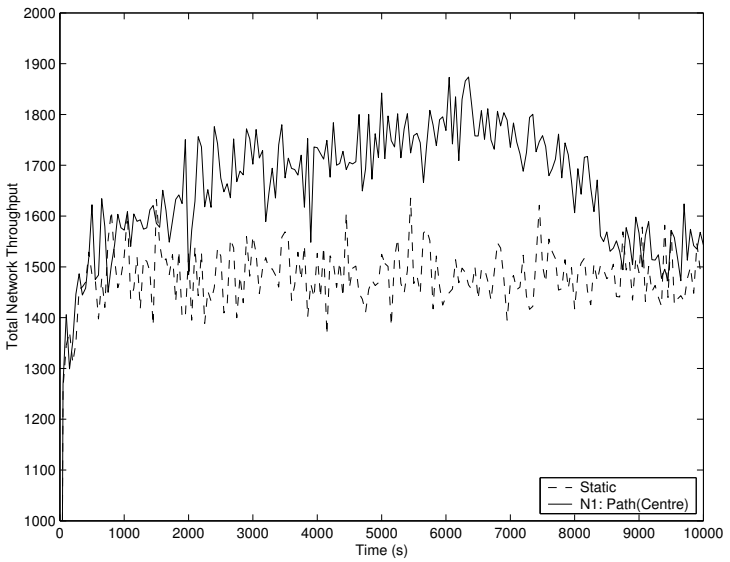

Figure 8: Comparison between the total throughput for the static network, and the networks with nodes $N1$ and $N8$ following mobility paths.

within the network *has* increased as $N1$ moves towards the centroid of the network. In comparison, when $N1$ moves from the centre to the edge of the network, the overall throughput decreases. Thus our results indicate ways in which the overall performance varies with the current geographical distribution of the users. Moreover, mobile users can affect not just their own performance, but also the overall performance of the network (see also [4]).

5 Conclusions

We have considered how incentives can be integrated into the operation of a mobile ad hoc network, so that the cost of resources consumed at transit nodes, when forwarding traffic along multi-hop routes, can be recovered using pricing mechanisms. These prices are determined in a distributed fashion, where algorithms are used by individual users to update their prices based on their bandwidth and power usage. Routes for connections from a user to a particular destination are chosen such that the route price is minimal. This forms a dual algorithm for traffic management within the network.

Incentives for collaboration have been provided through the concept of a user having a credit balance, which receives an initial endowment when the user joins the network. The credit balance accumulates notional credit accrued by forwarding traffic for other users, while any traffic generated from a particular user decreases the credit balance based on the cost of forwarding the traffic to its destination. The amount of traffic that a user can generate is directly related to its current credit balance—hence the user's incentive to both act as a transit node for other users and move to locations within the network where it can forward more traffic.

In this paper, we have studied this system through fluid-level simulations. These simulations have demonstrated that users' prices and credit balances stabilise for a static ad hoc network and shown the advantages in being near the centre of the network, as this allows nodes to act as transit nodes for a larger number of routes. We have also shown that mobility through the centre of the network can increase an individual user's throughput, as well as increasing the overall throughput of the system.

Further work includes exploring analytically the stability of the model, with the view of selecting appropriate parameters for updating user prices and discounting their credit balances. We will also investigate re-routing protocols that minimise the routing information that needs to be distributed in the network, while at the same time achieving near-minimal cost routing. In general, we have found that our model captures many of the fundamental trade-offs within the collaborative setting of an ad hoc network.

References

[1] L. Buttyan and J. P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *ACM/Kluwer Mobile Networks and Applications*, 8(5), October 2003.

[2] J.H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networks. *Proceedings of INFOCOM'00*, March 2000.

[3] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35:1969–1985, 1999.

[4] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10(4):477–486, August 2002.

[5] D. Johnson, D. Maltz, Y. C. Hu, and J. Jetcheva. The dynamic source routing protocol for mobile ad hoc networks (DSR). *IETF Internet Draft* `draft-ietf-manet-dsr-07.txt`, February 2002.

[6] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, March 1998.

[7] R. Liao, R. Wouhaybi, and A. Campbell. Incentive engineering in wireless LAN based access networks. *Proceedings of ICNP 2002*, November 2002.

[8] S. Low and D. Lapsley. Optimization flow control I: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, December 1999.

[9] P. Marbach and R. Berry. Downlink resource allocation and pricing for wireless networks. *Proceedings of INFOCOM'02*, June 2002.

[10] C. Perkins, E. Belding-Royer, and S. Das. Ad hoc on-demand distance vector (AODV) routing. *IETF Internet Draft* `draft-ietf-manet-aodv-12.txt`, November 2002.

[11] Vasilios Siris. Resource control for elastic traffic in CDMA networks. *Proceedings of MOBICOM'02*, September 2002.

[12] V. Srinivasan, C. Chiasserini, P. Nuggehalli, and R. Rao. Optimal rate allocation and traffic splits for energy efficient routing in ad hoc networks. *Proceedings of INFOCOM'02*, June 2002.

[13] W.H. Wang, M. Palaniswami, and S. Low. Optimal flow control and routing in multi-path networks. *Preprint submitted to Performance Evaluation*, October 2002.