

Coalition Games and Resource Allocation in Ad-Hoc Networks

R.J. Gibbens¹ and P.B. Key²

¹ Computer Laboratory, University of Cambridge, William Gates Building,
15 JJ Thomson Avenue, Cambridge, UK, CB3 0FD

`Richard.Gibbens@cl.cam.ac.uk`

² Microsoft Research Cambridge, Roger Needham Building,
7 JJ Thomson Avenue, Cambridge, UK, CB3 0FD

`Peter.Key@microsoft.com`

Abstract. In this paper we explore some of the connections between cooperative game theory and the utility maximization framework for routing and flow control in networks. Central to both approaches are the allocation of scarce resources between the various users of a network and the importance of discovering distributed mechanisms that work well. The specific setting of our study is ad-hoc networks where a game-theoretic approach is particularly appealing. We discuss the underlying motivation for the primal and dual algorithms that assign routes and flows within the network and coordinate resource usage between the users. Important features of this study are the stochastic nature of the traffic pattern offered to the network and the use of a dynamic scheme to vary a user's ability to send traffic. We briefly review coalition games defined by a characteristic function and the crucial notion of the Shapley value to allocate resources between players. We present a series of experiments with several test networks that illustrate how a distributed scheme of flow control and routing can in practice be aligned with the Shapley values which capture the influence or market power of individual users within the network.

1 Introduction

In this paper we explore some of the connections between cooperative game theory and the utility maximization framework for routing and flow control in networks. Central to both approaches are the allocation of scarce resources between the various users of a network and the importance of discovering distributed mechanisms that work well.

The specific setting of our study is ad-hoc networks and we examine the scheme proposed in [5].

The paper is organized as follows. In Sect. 2 we explain the basic model and quantities of interest largely following the notation used in [5]. We discuss the underlying motivation for the primal and dual algorithms that assign routes and flows within the network and coordinate resource usage between the users.

Important features of this study are the stochastic nature of the traffic pattern offered to the network and the use of a dynamic scheme to vary a user's ability to send traffic. We also review coalition games defined by a characteristic function and the crucial notion of the Shapley value to allocate resources between participants.

Section 3 presents a series of experiments with several test networks that illustrate how a distributed scheme of flow control and routing can in practice be aligned with the Shapley values which capture the influence or market power of individual users within the network.

2 Models

In this section we outline the basic models and quantities of interest. The essential features follow those given in [5].

2.1 Basic Models

Let N be the set of nodes and let \mathcal{R} be a set of routes. For each $j \in N$ write $\mathcal{R}^S(j) \subset \mathcal{R}$ for the set of routes which start at j and $\mathcal{R}^D(j) \subset \mathcal{R}$ for the set of routes which end at j .

For each source s we will denote by x_s the flow starting at s which flows at rate y_r on route $r \in \mathcal{R}^S(s)$ with

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

Then the amount of flow, c_j , through a node $j \in N$ is given by

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

where the first term aggregates all flows either starting or ending at j and where the second term aggregates all flows transiting both in and out of node j (and thus contribute twice to the quantity c_j). A node is constrained by some capacity, C_j , for aggregate flow so that

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

Similarly, we suppose that receiving and transmitting flows by a node consumes electrical power and we write γ_j for the power consumed at node j which we express in terms of the flows as

$$\gamma_j = \sum_{r \in \mathcal{R}^S(j)} y_r e_{jr}^T + \sum_{r \in \mathcal{R}^D(j)} y_r e^R + \sum_{r: j \in r \wedge r \notin \mathcal{R}^S(j) \cup \mathcal{R}^D(j)} y_r (e^R + e_{jr}^T). \tag{4}$$

Here we suppose that the power consumed at node j by a flow of rate y_r is $y_r e^R$ for receiving and $y_r e_{jr}^T$ for forwarding from j to the next node along the route r

after node j . We further suppose that receiving power consumption does not depend on the identity (and hence location) of the transmitter whereas transmitting does depend on the identity of the receiver. The total power consumed at node j is constrained by a quantity Γ_j , that is

$$\gamma_j \leq \Gamma_j \quad \forall j \in N. \quad (5)$$

2.2 Optimization Framework

Here we shall describe how flows x_s are determined given the constraints on bandwidth C_j and power consumption Γ_j .

We shall suppose the existence of prices μ_{jr} for use of node j by a unit amount of flow on route r and determine flows x_s and y_r by a *primal* algorithm such that

$$x_s = \sum_{r \in \mathcal{R}^S(s)} y_r = \frac{w_s}{\min_{r \in \mathcal{R}^S(s)} \sum_{j \in r} \mu_{jr}} \quad (6)$$

for given quantities w_s and with the proviso that y_r is only positive on routes r that attain the minimum in the denominator of the expression on the right-hand side. Thus, the action is to select flows such that the *rate of spending*, $x_s \sum_{j \in r} \mu_{jr}$, is minimal over the choice of routes $r \in \mathcal{R}^S(s)$ and has value w_s per unit time.

The underlying rationale for this primal algorithm is that of *proportional fairness* which adopts a maximization of utility with the specific choice $U(x) = w \log x$ as the utility function [4,7,10].

The prices μ_{jr} are intended to depend on current flows x_s in order to align demand for resources of bandwidth and power with their provision given in term of the quantities C_j and Γ_j . The dependence is through separate prices for bandwidth and power written μ_j^B and μ_j^P , respectively. Specifically, we write

$$\mu_{jr} = \begin{cases} e_{jr}^T \mu_j^P + \mu_j^B & j \text{ is the source for route } r \text{ (so } r \in \mathcal{R}^S(j)) \\ (e^R + e_{jr}^T) \mu_j^P + 2\mu_j^B & j \text{ is a transit node for route } r \\ & \text{(so } j \in r \wedge r \notin \mathcal{R}^S(j) \cup \mathcal{R}^D(j)) \\ e^R \mu_j^P + \mu_j^B & j \text{ is the destination for route } r \text{ (so } r \in \mathcal{R}^D(j)). \end{cases} \quad (7)$$

All the flows x_s, y_r and the various prices $\mu_{jr}, \mu_j^B, \mu_j^P$ will further depend on time and we use this dependence to specify the *dual* algorithm in which prices are adjusted over time according to the following equations

$$\frac{d}{dt} \mu_j^B(t) = \frac{\kappa \mu_j^B(t)}{C_j} (c_j - C_j) \quad (8)$$

$$\frac{d}{dt} \mu_j^P(t) = \frac{\kappa \mu_j^P(t)}{\Gamma_j} (\gamma_j - \Gamma_j) \quad (9)$$

where κ is a small positive constant.

2.3 Coalition Games and Shapley Values

We now turn to coalition games and their use in resource allocation problems. We suppose that a game is composed of a collection of N players corresponding to the source nodes in the ad-hoc network and that for each subset or coalition of players $S \subseteq N$ there is a payoff $v(S)$ given by the *characteristic function*. The characteristic function, $v : \mathcal{P}(N) \mapsto \mathbb{R}$, determines the maximum payoff that the coalition S can guarantee themselves by coordinating the actions of its members, whatever the other players decide. See [8,9] for further discussion of coalition games and the Shapley value approach.

We shall assume that $v(\emptyset) = 0$ and that $v(\cdot)$ is *superadditive*, that is

$$v(S \cup T) \geq v(S) + v(T) \quad (10)$$

whenever $S \cap T = \emptyset$.

An important notion for allocating the value $v(N)$ of the full coalition amongst the players is given by the *Shapley value* $\phi_i(v)$ defined by

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|! (|N| - |S| - 1)!}{|N|!} (v(S \cup \{i\}) - v(S)) . \quad (11)$$

It may be shown that the vector of Shapley values $(\phi_i(v) : i \in N)$ forms an *imputation*. That is, they are an assignment of the value $v(N)$ between the players where the assignment to player i is at least as great as they could obtain independently of the other players so that

$$\sum_{i \in N} \phi_i(v) = v(N) \quad (12)$$

$$\phi_i(v) \geq v(\{i\}), \quad \forall i \in N. \quad (13)$$

Note that to compute the Shapley values we require the characteristic function $v(S)$ to be determined for each of the 2^N possible coalitions of the full set of players N . In our experiments described later in Sect. 3 we have 10 players and so there are $2^{10} = 1024$ possible coalitions to consider.

The primal and dual algorithms are motivated by the underlying utility maximization framework and there is a large body of work that now supports that approach. A recent survey of this approach is given in [3]. The central notions are economic ones and relate to competitive Walrasian equilibria in exchange markets [8,10]. A connection between the non-cooperative notions of competitive equilibria and the cooperative notion of a Shapley value is through the *value equivalence theorem* of [2]. This work establishes how in a continuum setting of many small players the allocations associated with the Shapley values are the same as the competitive allocations. The approach we take studies these allocations in ad-hoc networks with a finite number of players.

The value equivalence theorem builds on earlier work that establishes a *core equivalence theorem* relating in a similar continuum setting the competitive allocations to those in the core of the game [1,8]. In general, the *core* of a cooperative game in characteristic form is the (possibly empty) set of imputations $(x_i : i \in N)$ such that

$$\sum_{i \in S} x_i \geq v(S) \quad \forall S \subset N. \quad (14)$$

We do not study the core further in this paper but instead concentrate on the use of the Shapley values.

2.4 Stochastic Features

The experiments discussed in Sect. 3 add stochastic features to the model discussed so far. We shall suppose that each source node s is controlled by a Markov process $D_s(t)$ that assigns the destination node for the current flow starting at s . We also allow $D_s(t)$ to take a further state, labelled 0, indicating that there is *no* current flow associated with source node s . Thus, $D_s(t) \in \{0, 1, \dots, N\} \setminus \{s\}$ and $D_s(t) = d$ ($d \neq 0$) means that the user s has flow starting at s and terminating at d . If $D_s(t) = 0$ then the user s is currently inactive and necessarily $x_s(t) = 0$. As the Markov processes $(D_s(t) : s \in N)$ change state then so do the sets $\mathcal{R}^S(\cdot)$ and $\mathcal{R}^D(\cdot)$ describing the sets of routes corresponding to the random source-destination pairs.

In our experiments the random holding times in the different states are independent exponentially distributed random variables with a common parameter λ . The permitted transitions of $D_s(t)$ are such that from the inactive state ($D_s(t) = 0$) the source will select any destination from $\{1, 2, \dots, N\} \setminus \{s\}$ equally likely. In the active state ($D_s(t) \neq 0$) the only transition is to the inactive state. Thus a source alternates between inactive and active periods with a destination node chosen uniformly at random for each successive active period.

The stochastic effects of the traffic patterns changing over time will accordingly imply a continual adjustment of flows and prices using the joint primal and dual algorithms. In our experiments we have further adopted the simplification that routes chosen (by the least cost primal algorithm) do not subsequently change during an active period even though prices may fluctuate to an extent that the chosen routes are no longer least cost ones.

Thus, the primal algorithm of equation (6) is revised to

$$x_s(t) = \frac{w_s(t)}{\min_{r \in \mathcal{R}^S(s)} \sum_{j \in r} \mu_{jr}(t)} \quad (15)$$

if $D_s(t) \neq 0$ and $x_s(t) = 0$ if $D_s(t) = 0$. The routes r for which $y_r > 0$ are determined by the minimum above when the flow initially becomes active and are then maintained without change throughout the active period.

2.5 Dynamic Schemes

A net balance of transit costs earned over those paid is maintained by the quantity, $b_s(t)$, defined for each $s \in N$ by

$$\frac{d}{dt}b_s(t) = \sum_{r \in \mathcal{R}^S(s)} y_r(t)\mu_{sr}(t) - w_s(t) \tag{16}$$

with the initial condition $b_s(0) = 1$. The first term measures revenue from transit fees per unit time and the second term, $w_s(t)$, as already noted, is the spending rate of source node s on its transit costs incurred by its flow of $x_s(t) = \sum_{r \in \mathcal{R}^S(s)} y_r(t)$. Here we allow the possibility that spending rates will vary over time as the function $w_s(t)$.

Furthermore, we assume that each source node has an initial endowment, $b_s(0)$, of one unit.

3 Experiments

Having reviewed the basic theory and choice of mechanisms which underly our model we now explore through a number of experiments the joint behaviour of the ad-hoc network.

Figure 1 shows the set of 10 nodes, here labelled $N = \{A, B, \dots, J\}$, placed uniformly at random in a square of side 100 units as considered in [5]. Edges are shown between pairs of distinct nodes that are a Euclidean distance of no more than 56 units apart. The purpose of these edges is to define the set of routes available for flow. A possible flow must be along a route corresponding to a path in the network. The network shown in Fig. 1 is connected but we shall also consider subnetworks defined by subsets of nodes (and just the edges incident to nodes within the subset). It is possible for such subnetworks to fail to be connected (the choice of distance threshold 56 controls which subnetworks are disconnected). In our experiments we have taken that for a disconnected subnetwork the flows are set to $x(t) = 0$ throughout the subnetwork. Other possibilities could include considering connected components separately.

We shall also consider in our experiments the effect of movement by the nodes to new locations. Figure 2 shows a new network where node B has moved from a location on the extreme of the network to a new location determined by the centroid of the remaining 9 nodes.

In the experiments we have taken the following choice of parameters. The gain parameter in the primal and dual algorithms is $\kappa = 0.05$, the power coefficients are $e^R = 0.001$ and $e^T = 0.0001 \times d$ where d is the Euclidean distance from the transmitting node to the receiver. The bandwidth capacity is $C_j = 10$ and the power constraint is $I_j = 0.5$. The distributions of the holding times for the active and inactive periods are independent exponential distributions with means $\lambda^{-1} = 0.5$ seconds.

In the next section we shall discuss the choice of the spending rate parameter, $w_s(t)$, for the primal algorithm and discuss two important classes of scheme: a static case and a dynamic case.




Fig. 1. A network of 10 nodes $N = \{A, B, \dots, J\}$ located uniformly at random in a square of side 100 units. An edge is shown between each pair of nodes separated by a Euclidean distance of at most 56 units.




Fig. 2. A second network of 10 nodes $N = \{A, B, \dots, J\}$: compared to the first network B has now moved to the centroid of the remaining 9 nodes




Fig. 3. Node balances over time with a static choice of $w_s(t) = 0.3$. The labels in the right hand margin identify the node(s). The node balances show clear trends: some increasing and some decreasing.

3.1 Static Schemes

We now describe a series of experiments with a static choice of spending rate parameter fixed over time at a value $w_s(t) = 0.3$ the same for each source node.

Figure 3 shows the net balance $b_s(t)$ over time for a simulation of duration 10,000 seconds. All the balances show either clear increasing trends or clear decreasing trends according to whether the spending rate of 0.3 units per second is above or below the rate of earning from transit fees charged to other flows.

Such imbalances between spending and earning reflect a disparity between the market power of the users and the allocation obtained. We demonstrate this by

Table 1. Sample means and standard deviations of node throughputs calculated from 50 independent replicates

Node	A	B	C	D	E	F	G	H	I	J
Mean	4.15	4.16	5.50	6.76	7.56	3.75	6.17	8.16	5.38	7.70
Standard deviation	0.02	0.02	0.03	0.03	0.05	0.02	0.04	0.05	0.02	0.06

considering the observed throughputs under the static scheme with the Shapley values constructed from a characteristic function formed from the system-wide throughput of the subcoalitions $S \subset N$ for each of the $2^{10} = 1024$ possible subcoalitions of the 10 nodes.

The specific details of the Shapley value calculation are described as follows. First, for each subcoalition, S , of users the subnetwork was tested for connectedness. If the subnetwork was disconnected then we set $v(S) = 0$. For the connected subnetworks we set

$$v(S) = \max_{S' \subseteq S} X(S') \tag{17}$$

where $X(S')$ was the observed system-wide throughput of subnetwork S' . Note that $X(S') = 0$ if S' is disconnected. The use of the maximum over all subcoalitions S' was to ensure the superadditivity property of the characteristic function. In some subnetworks, S , it was observed that $X(S') > X(S)$ for $S' \subset S$ which prevents taking $v(S) = X(S)$ for the characteristic function.

The definition in equation (17) thus admits the coordinated action of the players of the subcoalition to drop a player if that would strictly *increase* system-wide throughput even if this wasn't the observed behaviour of the scheme when simulated. Further performance metrics besides system-wide throughput could easily be incorporated into the definition of the quantity $X(\cdot)$. See [6] for an alternative means of ensuring the superadditivity property which has important connections with the game-theoretic notion of a *stable set* of imputations.

The random quantity, $X(S)$, was estimated in our experiments by a long-run average over the randomly varying traffic patterns driven by the Markov processes $(D_s(t) : s \in N)$. Table 1 shows the sample means and sample standard deviations of the node throughputs from 50 independent simulation replicates. The standard deviations show little variability in the estimates of the mean.

Figures 4 and 5 show the correspondence between the observed throughputs and the Shapley values and in each case the proportion of throughput or of the sum of the Shapley values, $v(N) = \sum_{i \in N} \phi_i(v)$ is given. As just noted it is possible for the system-wide throughput to be less than $v(N)$ and so we consider proportions only throughout our comparisons.

Figure 4 shows the Shapley values and observed throughputs according to the locations of the nodes. It is clear that nodes at extreme locations (such as B, C, E and H) receive far larger shares of the system-wide throughput than is allocated by the share of the Shapley value. Conversely, nodes close to the centre




Fig. 4. Shapley values and observed values for the proportion of throughputs by node with static choice of $w_s(t) = 0.3$. The radius of the circles measure either the proportion of the Shapley value or the observed throughput.

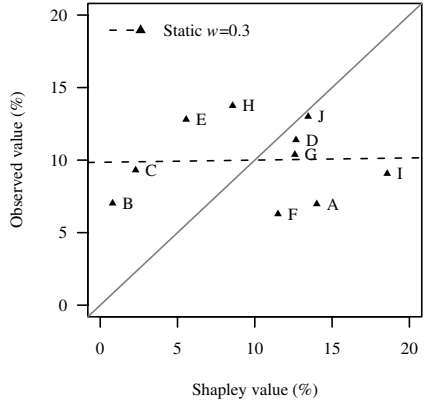


Fig. 5. Scatter plot of observed values and Shapley values with the static choice of $w_s(t) = 0.3$. The near horizontal dashed line is a least squares fit to the data points.

of the network (such as A, F and I) receive smaller shares of the system-wide throughput than those allocated according to the Shapley value.

These effects are also apparent from Fig. 5 which shows a scatter plot of the shares of the throughputs and Shapley values. Also shown here is the dashed line given by least squares fit to the data which is far from diagonal.

The dynamic models of the next set of experiments attempt to correct this bias which favours nodes at the extreme of the network with little market power in preference to those near the centre of the network with the largest Shapley values.

3.2 Dynamic Schemes

Here we take the spending rate parameters as the functions of the balance $b_s(t)$ given by

$$w_s(t) = \alpha b_s(t) \tag{18}$$

for a constant $\alpha \in (0, 1)$. For the experiments discussed here we took $\alpha = 0.3$. In this way a larger balance feeds through to a higher spending rate and if the net balance drops the spending rate is reduced accordingly.

Figure 6 shows the net balances under the dynamic scheme. All balances start at $b_s(0) = 1$ given by the initial endowment and then evolve over time to fluctuate about constant levels. The sum of the net balances, $\sum_{s \in N} b_s(t) = N = 10$ remains fixed over time since all spending by one node is earned by other nodes.




Fig. 6. Node balances over time with a dynamic choice of $w_s(t) = \alpha b_s(t)$. The labels in the right hand margin identify the node(s). Here node balances fluctuate about values without any long-term trend to increase or decrease. The lines shown are smoothed according to a moving average of 250 seconds.




Fig. 7. Shapley values and observed values for proportion of throughputs by node with dynamic choice of $w_s(t) = \alpha b_s(t)$. There is a much closer correspondence between the Shapley values and the observed throughputs in the dynamic case.




Fig. 8. Scatter plot of observed values and Shapley values, including static and dynamic choices of $w_s(t) = \alpha b_s(t)$. The near diagonal dashed line is a least squares fit between the Shapley values and the observed throughputs in the dynamic case.

From Fig. 6 we can see that nodes B and C have net balances $b_s(t)$ that converge to zero while node I has the highest net balance. Figures 7 and 8 show the shares of system-wide throughput and of Shapley value obtained under the




Fig. 9. Shapley values and observed values after node B has moved to the centroid




Fig. 10. Scatter plot of the proportion of Shapley values and observed throughputs after node B has moved to the centroid

Table 2. Estimated node throughputs under the dynamic scheme in the two networks

Node	A	B	C	D	E	F	G	H	I	J
B at extreme	5.81	0.01	0.02	9.06	2.27	4.81	8.25	3.45	8.76	10.02
B at centroid	5.86	11.16	0.02	10.51	2.19	3.52	9.05	3.28	11.56	9.92

dynamic scheme. We can see that there is a much closer correspondence between the observed share of the throughput obtained by each player and that given by the Shapley value approach. Thus market power and outcomes have been more closely aligned than under the static scheme. The line of least squares fit to the dynamic data is very close to the diagonal line and the departures from the diagonal line are more modest than for those in the static scheme.

Figures 9 and 10 show the dynamic scheme in operation in the second network where node B is no longer on the periphery of the network but placed at the centroid of the remaining 9 nodes. The figures show that the shares of the system-wide throughput and of the Shapley value are quite tightly aligned around the diagonal line in Fig. 10. In [5], the authors considered a trajectory for node B which started at the extreme position and passed through the centroid position and then on towards the upper boundary of the square. The system-wide throughput increases as B moves along this trajectory towards the centroid. In our experiments the throughput increased from 52.45 to 67.07. In both network scenarios the allocations obtained by the dynamic scheme closely align with Shapley values and thus the market power of the node. Table 2 shows the

node throughputs obtained under the dynamic scheme in the two networks and reveals that the majority of the additional system benefit arising as B moves to the more favourable position at the centroid accrues to B itself.

4 Conclusions

In this paper we have studied the scheme for resource allocation given in [5] for ad-hoc networks and have explored the connections between this scheme and notions from cooperative game theory. The Shapley value is one such notion that has enabled a broader understanding of how resource allocation takes place with this scheme.

Further work remains to investigate how widely these connections between cooperative game theory and the underlying utility maximization framework for flow control and routing in networks can be extended.

Acknowledgements. RJG acknowledges support from the UK EPSRC grant reference GR/S86266/01 and the International Technology Alliance in Network and Information Science (ITA).

References

1. Aumann, R.J.: Markets with a continuum of traders. *Econometrica* 32, 39–50 (1964)
2. Aumann, R.J.: Values of markets with a continuum of traders. *Econometrica* 43(4), 611–646 (1975)
3. Chiang, M., Low, S.H., Calderbank, A.R., Doyle, J.C.: Layering as optimization decomposition: A mathematical theory of network architectures. *Proceedings of the IEEE* 95(1), 255–312 (2007)
4. Chiang, M.: Balancing transport and physical layers in wireless multihop networks: Jointly optimal congestion control and power control. *IEEE Journal on Selected Areas in Communications* 23(1), 104–116 (2005)
5. Crowcroft, J., Gibbens, R.J., Kelly, F.P., Östling, S.: Modelling incentives for collaboration in mobile ad hoc networks. *Performance Evaluation* 57, 427–439 (2004)
6. Gillies, D.B.: Solutions to general non-zero-sum games. In: Tucker, A.W., Luce, D.R. (eds.) *Contributions to the Theory of Games*, vol. IV. Princeton University Press, Princeton (1959)
7. Kelly, F.P., Maulloo, A.K., Tan, D.K.H.: Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* 49, 237–252 (1998)
8. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Microeconomic Theory*. Oxford University Press, Oxford (1995)
9. Thomas, L.C.: *Games, Theory and Applications*. Dover Publications (2003)
10. Varian, H.R.: *Microeconomic Analysis*, 3rd edn. W.W. Norton and Company (1992)