STRATEGIC NEGOTIATION OF BANDWIDTH IN COOPERATIVE NETWORKS

Jonathan Bredin
Colorado College
Department of Mathematics, 14 E Cache la Poudre St, Colorado Springs, CO 80903, USA

Keywords: Cooperative networks, resource allocation, game theory, self-interested agents

Abstract: We analyze the scenario where a pair of network devices each periodically relies on the other to handle network traffic. Without immediate reward, the forwarding device incurs an opportunity cost in handling the other's request. We find, however, situations where rational decision makers prefer bandwidth exchange to isolated operation. We base our analysis on a take-or-leave-it protocol inspired by the Rubinstein bargaining model, and extend it to evaluate repeated interaction between pairs of devices.

1 INTRODUCTION

Around the world, wireless network users consolidate their resources to form cooperative networks to extend the range of their wirelessly networked devices (Flickenger, 2003). Devices in each of these networks are owned by multiple self-interested individuals and each device depends on its peers to handle some of its network traffic. The deployment cost for a cooperative network spreads evenly among the many users. Many wireless cooperative networks have emerged in urban neighborhoods around the world. Furthermore, the distributed ownership and policy development make cooperative networks potentially dynamic and robust to fluctuating user demand and mobile nodes.

The network depends device owners devices to *choose* to cooperate– there is no other motivation for participation in the network. Most cooperative networking organizations assume that users willingly donate hardware to build the network, but we find that there are several reasons to question the assumption. First, network applications require increasingly greater bandwidth, especially in light of ubiquitous multimedia, and bandwidth will become constrained. Second, in the wake of recent denial-of-service attacks from email viruses and worms, it is clear that a networked commons could benefit from some regulation. Finally, many obvious applications of cooperative networks exist in mobile wireless networks, where devices face not only bandwidth, but also battery power constraints, so network participation may not be rational.

This paper examines the motivations for peer-wise bandwidth exchange in cooperative networks. We find that even a short-sighted analysis provides incentive for rational decision makers to contribute to cooperative networks. The equilibrium reached in the myopic analysis, however, is not stable to perturbation or incomplete information.

In section 2 we present our model for bandwidth exchange. Section 3 describes the method that we use to compare the bandwidth allocations. From the evaluation and ideas from existing negotiation models, we derive a policy for one device to rationally decide whether or not to share bandwidth with another on the expectation of future cooperation. We describe some relevant work in section 6 and conclude in section 7 with directions to extend our analysis.

2 BANDWIDTH EXCHANGE IN COOPERATIVE NETWORKS

That the cost of deployment falls evenly among the network users motivates the construction of cooperative networks. Users contribute devices to form the network. Sometimes a device will link the cooperative network to the rest of the Internet. Frequently, the devices comprising a cooperative network are wirelessly linked using 802.11 protocols where most de-

Bredin J. (2004).
STRATEGIC NEGOTIATION OF BANDWIDTH IN COOPERATIVE NETWORKS.
In *Proceedings of the Sixth International Conference on Enterprise Information Systems*, pages 230-235
DOI: 10.5220/0002627302300235

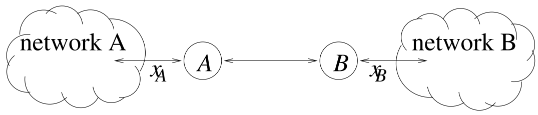

Figure 1: The network topology to be modeled.

vices could not individually connect to the Internet.

Several possible factors motivate Internet uplink sharing. Distributing access of several uplinks among network users spreads the load evenly among the uplinks. Users benefit from statistical multiplexing: the chance that all users require network access is small and the residual network capacity can be divided among the active users. The scenario also provides redundancy when one user's service provider fails to provide Internet access due to link or router failure, network overload, etc.

Another motivation may stem from that while cooperating network members may be mobile, their uplinks are statically located. A user that strays far from her Internet uplink may wish to use another uplink during her travel. The user "pays" for her access through providing other members with network access when they travel near her uplink.

Finally, financial reasons may motivate sharing of a network uplink. Perhaps a group of users lives in a building ill suited for network access; wiring the entire building may be expensive or intrusive and users could instead share a set of uplinks. In rural environments, a group of users may share a satellite uplink that would be too expensive for a single user to maintain. The shared payment of the Internet service may be arranged out of band, but to function, there must be incentive for devices to contribute to the rest of the network rather than freeload on the efforts of others.

This altruistic behavior, or at least the incentive to behave as such, is a central obstacle for cooperative-network operation. Forwarding network traffic incurs an opportunity cost of network bandwidth usage on the part of altruistic devices, i.e. the more one device helps another, the less work it can perform for itself. Rationally, donations come with some expectation that favors will be returned.

We model a scenario where two devices, A and B, repeatedly query one another to forward network traffic. A link partitioning the network connects the two devices. Device i connects to its side of the network with an x_i bit per-second link. Figure 1 sketches the network topology. For now, we ignore the bandwidth constraint of the connecting link as well as battery limits.

Periodically, device j requests that i give it network access. We assume that a device will not make another request until the current request is satisfied. Device j informs i the desired duration of usage; and i immediately returns an answer of x_{ij}, the amount of bandwidth it will provide to j.[1] If i denies j's request, no further exchange will occur between the devices. We assume that there is no cost in dividing bandwidth, so that

$$x_i = x_{ii} + x_{ij}. \quad (1)$$

3 UTILITY MODEL

We now assert some general assumptions we make that will drive the derivation of sharing policy in the next section. Device i generates utility at a rate of $u_{ij}(x)$ per second, where u_{ij} is a continuous function of x, the amount of j's bandwidth available for i to consume. More bandwidth is always better, but not proportionally so. Hence du_{ij}/dx is positive and d^2u_{ij}/dx^2 is negative. Without loss of generality, we root the utility function so that $u_{ij}(0) = 0$. Furthermore, u_{ij} is onto with respect to the non-negative reals, modeling the belief that enough bandwidth will solve any problem. We will use $u_{ij}(x) = \log(1+x)$ for the graphical and numeric examples for the remainder of the paper.

We describe device i's total utility as the sum of two parts: the utility derived from local bandwidth usage, and utility derived from consuming another device's bandwidth. We write the expected utility function as

$$U_i(x_{ii}, x_{ji}) = \alpha_i \hat{U}_{ii}(x_{ii}) + (1-\alpha_i)\hat{U}_{ij}(x_{ji}), \quad (2)$$

where α_i is a weight in $[0,1]$ to determine the importance of local versus remote consumption. The functions $\hat{U}_{ii}$ and $\hat{U}_{ij}$ represent the expected utility derived from local and remote bandwidth, respectively. Since bandwidth usage varies over time, we calculate the time-discounted expected utility from each component as a function of the estimated bandwidth consumption presented in the next section.

3.1 Time-discounted evaluation

Devices prefer consuming earlier than later. We define the value of i's consumption of j's bandwidth from time t_1 until t_2 as

$$U_{ij}(x, t_1, t_2) = \int_{t_1}^{t_2} u_{ij}(x) e^{-t\gamma_i}\, dt = \left. \frac{u_{ij}(x)e^{-t\gamma_i}}{\gamma_i} \right|_{t_2}^{t_1}, \quad (3)$$

[1] For rest of the paper, the first subscript of a bandwidth amount denotes the provider, whereas the second, possibly identical, subscript denotes the consumer.

where $\gamma_i \in [0,1]$ is i's discount factor.

We now calculate the expected time-discounted utility a device generates from consuming bandwidth. For its own bandwidth, a device i values its choice to consume x_{ii} bandwidth when it is requested, and x_i at other times, as

$$\hat{U}_{ii}(\hat{x}_{ii}) = \int_{r=0}^{\infty}\int_{t=0}^{\infty} q_j(r)p_j(t)\Big[U_{ii}(x_i,0,t) + U_{ii}(\hat{x}_{ii},t,t+r) + \hat{U}_{ii}(\hat{x}_{ii})e^{-\gamma_i(t+r)}\Big]\,dt\,dr, \quad (4)$$

where we denote the probability density function for the duration that j requests bandwidth from i as $q_j(r)$, the inter-request time probability density function as $p_j(t)$, and the estimate i has of using its own bandwidth as $\hat{x}_{ii}$.

To simplify the calculation, we assume that each device generates requests at a Poisson rate with exponentially distributed request durations, though our decision model in section 4 does not depend on the assumption. Under the Poisson assumption where device i has a request intensity of λ_i and an average request duration of $1/\phi_i$, the time-discounted expected value of a device consuming its own bandwidth is

$$\hat{U}_{ii}(\hat{x}_{ii}) = \frac{(\phi_j+\gamma_i)u_{ii}(x_i) + \lambda_j u_{ii}(\hat{x}_{ii})}{\gamma_i(\gamma_i+\phi_j+\lambda_j)}. \quad (5)$$

The calculation of the value device i has for consuming j's bandwidth is a similar,

$$\hat{U}_{ij}(\hat{x}_{ji}) = \int_{r=0}^{\infty}\int_{t=0}^{\infty} q_i(r)p_i(t)\Big[U_{ij}(\hat{x}_{ji},t,t+r) + \hat{U}_{ij}(\hat{x}_{ji})e^{-\gamma_i(t+r)}\Big]\,dt\,dr, \quad (6)$$

and the Poisson assumption simplifies to the result

$$\hat{U}_{ij}(\hat{x}_{ji}) = \frac{\lambda_i u_{ij}(\hat{x}_{ji})}{\gamma_i(\gamma_i+\phi_i+\lambda_i)}. \quad (7)$$

We note that, even without the Poisson assumption, the utility of consuming no bandwidth is still zero.[2]

3.2 Feasible allocations

For a cooperative network to function, a device must be better off periodically sharing than operating in isolation and consuming only it its own bandwidth. In this section we model a device's preference for consuming its own bandwidth in comparison to its partner's. We use the comparison to derive the set of rationally feasible bandwidth allocations from which the devices may agree.

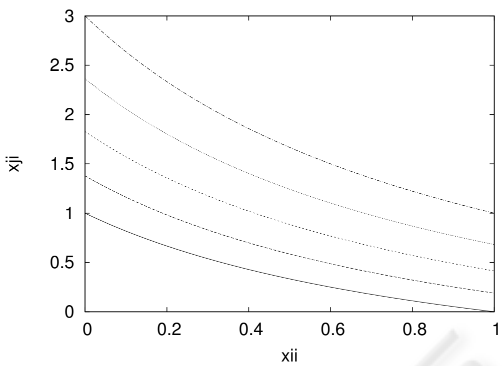

Figure 2: Five isoquants of i's bandwidth and j's bandwidth donated to i. The lowest curve represents the set of allocations for which i is indifferent to sharing or acting on its own. The highest curve represents a set of allocations that yield twice as much utility to i than operating in isolation.

Let us assume that devices attempt to maximize a weighted sum of their valuations for bandwidth usage,

$$U_i(\hat{x}_{ij},\hat{x}_{ji}) = \alpha_i\hat{U}_{ii}(x_i-\hat{x}_{ij}) + (1-\alpha_i)\hat{U}_{ij}(\hat{x}_{ji}), \quad (8)$$

where $\alpha_i \in [0,1]$ expresses i's preference for local versus remote bandwidth usage.

An isoquant is a set of equally preferable allocations. As a result of modeling utility as an increasing function of bandwidth with diminishing returns, the curves representing isoquants of two commodities will always be convex.(Mas-Colell et al., 1995) We plot some isoquants for a device's preference for bandwidth in figure 2. The lowest curve represents the set of allocations that are all equally preferable to operating without cooperation between the two devices. It would not be rational for a device to operate with the expectation of receiving an allocation lying under the lowest curve.

Since we assume that dividing bandwidth is costless, we can rearrange the isoquant curves in figure 2 in terms of $x_{ij} = x_i - x_{ii}$ and x_{ji}. We can also do the same for j's indifference curves and plot the two devices' isoquants against each other. Figure 3 plots each device's isoquant curves at the level of utility indifferent to isolated operation. The figure demonstrates the set of feasible bandwidth allocations.

The lens-shaped region between the two curves represents a set of allocations with which both devices prefer sharing to isolated operation. With any

[2]Contact the author for a full version of the paper with an appendix that shows that both $\hat{U}_{ii}$ and $\hat{U}_{ij}$ are increasing-concave in the utility consumed.

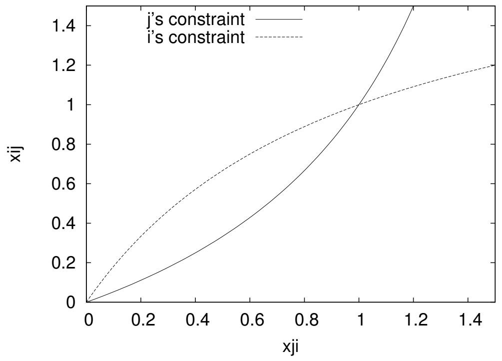

Figure 3: The set of feasible allocations is the lens shaped area between the curves. Device i will opt to consume only its own bandwidth if the allocation falls on the left side of its constraint curve, whereas j will only consume above its constraint curve.

other allocation, one device would prefer to opt out of cooperation. We use this area to derive a rational bandwidth-sharing policy in the next section.

4 NEGOTIATION

We now use the utility model to derive a rational policy to share bandwidth with another device. The devices repeatedly ask each other for assistance. After several exchanges, device i has an expectation of receiving $\hat{x}_{ji}$ bandwidth and giving $\hat{x}_{ij}$ when necessary.

We use a rationale taken from the Rubinstein bargaining model: in bargaining, an agent offers its partner an allocation whose outcome the partner is indifferent to accepting or not (Rubinstein, 1982). If we assume that the impact of the next interaction is negligible– perhaps the devices are mobile and will move on anyway– then a device will offer its partner an allocation that the other device prefers equally to solo operation. The allocation satisfies

$$U_j(\hat{x}_{ji}, x_{ij}) = U_j(x_j, 0). \tag{9}$$

Figure 3 sketches with a solid line the offer that i makes to j as a function of j's generosity to satisfy equation 9. We can solve for the the lower boundary, device i's response as a function of j's cooperation, as

$$x_{ij} = h_i(\hat{x}_{ji}) = \hat{U}_{ji}^{-1}\left(\frac{\alpha_j\left(u_{jj}(x_j) - \gamma_j \hat{U}_{jj}(x_j - \hat{x}_{ji})\right)}{(1-\alpha_j)\gamma_j}\right). \tag{10}$$

In the example where utility generation is a logarithmic function of bandwidth, we derive the boundary

$$x_{ij} = h_i(\hat{x}_{ji}) = u_{ji}^{-1}(k_{ji}(u_{jj}(x_j) - u_{jj}(x_j - \hat{x}_{ji}))), \tag{11}$$

where

$$k_{ji} = \frac{\alpha_j \lambda_i(\gamma_j + \phi_j + \lambda_j)}{(1-\alpha_j)\lambda_j(\gamma_j + \phi_i + \lambda_i)}. \tag{12}$$

With the rationality constraint that i may prefer to share less than the amount required to sustain trade with j, and i reverts back to its own constraint. At this extreme case, i's response is simply the inverse of j's initial response. We incorporate the two cases to yield i's response g_i,

$$x_{ij} = g_i(\hat{x}_{ji}) = \begin{cases} h_i(\hat{x}_{ji}) & U_i(h_i(\hat{x}_{ji}), \hat{x}_{ji}) \geq U_i(0, x_i) \\ h_j^{-1}(x_{ji}) & U_i(h_i(\hat{x}_{ji}), \hat{x}_{ji}) < U_i(0, x_i) \end{cases}. \tag{13}$$

We can visualize i's response as the lower boundary of the two curves plotted in figure 3.

5 EQUILIBRIUM

It is useful not only to know how devices immediately share bandwidth, but also whether they will continue to share bandwidth. To this end, we study equilibrium behavior to derive conditions under which devices will exchange bandwidth. Furthermore, equilibria stability also interests us; we wish to know how small changes in exchange affect exchange decisions.

We describe equilibrium as the pair of responses such that the equality

$$x_{ij} = g_i(g_j(x_{ij})) \tag{14}$$

holds. Notice that with the response derived in the previous subsection, the trivial equilibrium at $x_{ji} = x_{ij} = 0$ is always feasible, but sometimes other equilibria exists. We treat the composition from equation 14 as an iterated function system to determine equilibria and their stability. A necessary and sufficient condition for stability of an equilibrium is that the absolute slope of the iterated function system is less than one (Kaplan and Glass, 1995). With our scenario, the condition becomes

$$1 > \left|\frac{dg_i(g_j(x_{ji}))}{dx_{ji}}\right| = \left|\frac{dg_i}{dx_{ji}}\frac{dg_j}{x_{ij}}\right|. \tag{15}$$

The intuition is that when the absolute slope is less than one, changes in a device's output cause the other device to respond, but to a lesser extent. An absolute slope greater than one means that one device's change is magnified by its partner, and the equilibrium is unstable.

We rewrite a device's response function as

$$g_i(x_{ji}) = \min\{h_i(x_{ji}), h_j^{-1}(x_{ji})\}. \quad (16)$$

Note that because h_i is continuous and convex, there is a second equilibrium where h_i and h_j^{-1} intersect beyond $x_{ji} = 0$. Additionally, i's partner shares the same intersection. The allocation represented at the allocation yields utility for each device equal to not sharing any bandwidth, so the rationality constraint is enforced for each device. Let us name the equilibrium where $h_i = h_j^{-1}$ as (x_{ji}^*, x_{ij}^*) and define

$$k = x_{ij}^*/x_{ji}^*. \quad (17)$$

We now will use k to show that allocations to the right of the second equilibrium will result in allocations closer to to it, and allocations to the left of the equilibrium will result in devices eventually boycotting trade.

Theorem 1 *The equilibrium* (x_{ji}^*, x_{ij}^*) *attracts allocations* $(x_{ji}, g_i(x_{ji}))$ *whenever* $x_{ji} > x_{ji}^*$.

Proof. We prove the theorem by showing that the slope of the iterated response $g_i(g_j(x_{ij}))$ is strictly less than one. Since each device's rationality constraint is increasing concave in the amount it receives from its partner, k provides a boundary on the slope of each device's response. The ratio k provides a strict upper bound on i's marginal response, and $1/k$ is the strict upper bound for j's marginal response. Additionally, both these marginals are positive, since the constraints are strictly increasing. Hence for $x_{ji} > x_{ji}^*$

$$\frac{dg_i\dot{g}_j}{dx_{ji}} = \frac{dg_i}{dx_{ji}}\frac{dg_j}{x_{ij}} < k * 1/k = 1, \quad (18)$$

so the equilibrium (x_{ji}^*, x_{ij}^*) is attractive from the right since the derivative of the iterated response is less than one, but greater than zero. $\diamond$

Theorem 2 *The equilibrium* (x_{ji}^*, x_{ij}^*) *repulses allocations* $(x_{ji}, g_i(x_{ji}))$ *whenever* $x_{ji} < x_{ji}^*$.

Proof. We use a similar argument to show that the marginal iterated response is greater than one when x_{ji} is less than x_{ji}^*. Note that the devices' response functions are increasing convex for the values of interest and that $g_i(0) = 0$. Again, k provides a bound for marginal response near the equilibrium $(x_{ji}^*, g_i(x_{ji}^*))$. This time, $dg_i/dx_{ji} > k$ and $dg_j/dx_{ij} > 1/k$, so

$$\frac{dg_i\dot{g}_j}{dx_{ji}} = \frac{dg_i}{dx_{ji}}\frac{dg_j}{x_{ij}} > k * 1/k = 1. \quad (19)$$

Hence, from the left, the equilibrium $(x_{ji}^*, g_i(x_{ji}^*))$ is repulsive. $\diamond$

While there is an equilibrium that involves both devices exchanging bandwidth, the devices are indifferent to the equilibrium and not sharing. Furthermore, the equilibrium is unstable in the sense that deviations in exchange always yield less trade.

6 RELATED WORK

Rubinstein presents the seminal bargaining model where two agents take turns proposing allocations to each other (Rubinstein, 1982). When the total value to be allocated decays over time, the scenario ends with a finite number of iterations. Kraus applies the Rubinstein model to solve automated service negotiation problems (Kraus, 2001). Muthoo studies situations where agents repeatedly enter negotiations with each other (Muthoo, 1999). He only considers scenarios where agents alternate taking turns in initiating negotiation, and finds that when when agents have the opportunity to make only one offer, as in our model, that devices receive all or nothing.

The idea that networks can benefit from cooperative behavior is not new to this paper. Padmanabhan and Sripanidkulchai analyze network traces of news-flash traffic to create routing policies where devices cooperate to reduce latency of network broadcasts (Padmanabhan and Sripanidkulchai, 2002).

Buttyan and Hubaux observe that devices may not wish to cooperate, but noncooperation eliminates many applications (Buttyan and Hubaux, 2003). They explore the possibility of embedding trusted hardware into every network device. The trusted hardware enforces that each device may not use the network for its own purposes more than it forwards other devices' messages. They present several policies to optimize local performance under the management of the censoring trusted hardware.

Recently, several research groups have addressed incentive problems for participation in peer-to-peer file-sharing applications, focusing on modeling incentives through repeated multi-player prisoners' dilemma games. Lai et al. measure the value of private versus shared reputation information (Lai et al., 2003). Their model addresses "whitewashing," the possibility that a device can change its identity for each network interaction, and shows that systems that rely on private reputation information do not scale well. Ranganathan et al. compare reputation and market-based incentive mechanisms to find scenarios where reputation-based model performance approaches the performance of market-based models (Ranganathan et al., 2003).

7 DISCUSSION

We present a pessimistic model for network exchange to show that there exist equilibria allocations where

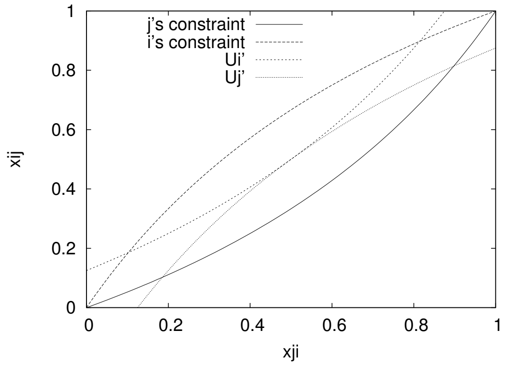

Figure 4: The set of feasible allocations and device i and j's isoquants that include the allocation that maximizes social surplus.

rational devices choose to cooperate with one another. The model is pessimistic in that each device myopically assumes that its actions will not affect its partner's future decision, except to wholesale discontinue trading.

We note that there are frequently better allocations than the ones that our model arrives. Figure 4 sketches the set of feasible allocations as well as two new isoquants for the devices that maximize the sum of the devices' utilities– the social surplus.

We have two reasons to have optimism regarding more efficient allocation. First, more-forward looking models will likely yield allocations lying deeper into the lens-shaped region. We are currently investigating negotiation models where each device reasons about the effects its own offer has in driving the pair of devices towards a favorable equilibrium allocation. Muthoo discovers that agents that alternate positions in repeated bargaining situations strive for more socially optimal allocations whereas in the traditional single-shot bargaining game, a device receives more resources the less it values the resource as a result of brinksmanship. In Muthoo's repeated bargaining model, the reverse, more socially efficient, effect manifests (Muthoo, 1999).

The second source of intuition seems paradoxical: competition will drive allocations to which devices strongly prefer over not sharing. The idea is that given the choice of sharing with two devices, a rational decision maker will choose the one that yields a more desirable allocation. So far, we have only modeled bilateral bandwidth exchange. That competition might provide sweeter allocations gives us further incentive to continue our research. It is likely that testing multilateral exchange will be beyond scope of an analytic model, so we are currently implementing our bandwidth allocation policies inside the ns network simulator.

That even a pessimistic model can motivate devices to exchange bandwidth motivates us to further study incentives for participation in cooperative networks. We are excited to investigate foresighted models of exchange as well as multilateral exchange to motivate more efficient exchange and to study larger networks.

REFERENCES

Buttyan, L. and Hubaux, J.-P. (2003). Stimulating cooperation in self-organizing mobile ad hoc networks. *Mobile Networks and Applications*, 8(5). to appear.

Flickenger, R. (2003). *Building Wireless Community Networks, 2nd Edition*. O'Reilly, Sebastopol, CA.

Kaplan, D. and Glass, L., editors (1995). *Understanding Nonlinear Dynamics*. Springer Verlag, New York, NY.

Kraus, S. (2001). *Strategic Negotiation in Multiagent Environments*. MIT Press, Cambridge, MA.

Lai, K., Feldman, M., Stoica, I., and Chuang, J. (2003). Incentives for cooperation in peer-to-peer networks. In *Proceedings of the First Workshop on Economics of Peer-to-Peer Systems*.

Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press, New York, NY.

Muthoo, A. (1999). *Bargaining Theory with Applications*. Cambridge University Press, Cambridge, UK.

Padmanabhan, V. K. and Sripanidkulchai, K. (2002). The case for cooperative networking. In Druschel, P., , Kaashoek, F., and Rowstron, A., editors, *Peer-to-Peer Systems*, pages 178–190. Springer-Verlag, New York, NY.

Ranganathan, K., Ripeanu, M., Sarin, A., and Foster, I. (2003). To share or not to share: An analysis of incentives to contribute in collaborative file sharing environments. In *Proceedings of the First Workshop on Economics of Peer-to-Peer Systems*.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109.