Smoothing Out Focused Demand for Network Resources

Kevin Leyton-Brown, Ryan Porter, Shobha Venkataraman, Balaji Prabhakar
CS Department, Stanford University, Stanford CA 94305
{kevinlb;rwporter;shobhav;balaji}@cs.stanford.edu

ABSTRACT

We explore the problem of sharing network resources when agents' preferences lead to temporally concentrated, inefficient use of the network. In such cases, external incentives must be supplied to smooth out demand. Taking a game-theoretic approach, we consider a setting in which bandwidth is available during several time slots at a fixed cost, but all agents have a natural preference for choosing the same slot. We present four mechanisms that motivate agents to distribute load optimally by probabilistically waiving the cost for each time slot, and analyze equilibria.

1. INTRODUCTION

It is common for networks to experience frequent congestion even when average demand for the network is much less than the network's capacity. In some networks, times of peak demand are regular and predictable. Such *focused loading* can occur because many agents' utility functions are maximized by using the network at some focal time. For example, studies of long-distance telephone networks show a spike in usage when rates drop in the evening [7, 1]. Predictably heavy loads also occur on web servers just before deadlines or just after new content or services are made available. In this paper, we provide a game-theoretic analysis of several solutions to the problem of focused loading.

There exists a substantial body of existing work on managing congestion in networks. In particular, the problem of designing congestion control and pricing mechanisms to provide differentiated qualities-of-service (QoS) in the Internet has received a lot of attention. The essential issue is allocating network bandwidth fairly among concurrent users, given that agents are likely to act selfishly to maximize the bandwidth available to them [9]. This problem can be addressed with new technology: the network can isolate packet flows by erecting "bandwidth firewalls" to ensure fairness or approximate fairness [3, 4]. An alternate line of research takes an economic approach to congestion management. The network attempts to induce agents to condition their flows to prescribed parameters, avoiding the implementation complexity inherent in the technological approach [6, 5, 8].

Separate consideration of the case of focused loads is worthwhile for two main reasons. First, focused loading occurs at very predictable times. This means that it is possible to know in advance the cases for which such a specialized solution should be used. Second, specialized solutions can do a better job of dealing with the problem of focused loading than more general approaches. Focused loading occurs because agents have similar utility functions—particularly, functions that are maximized by using the network at a particular time. General congestion management techniques cannot take this information into account; however, additional knowledge about agent utility functions makes it possible to design mechanisms that collect more revenue and make fewer (e.g., computational) demands on the network.

2. PROBLEM DEFINITION

We begin with a canonical example. Consider a telephone network in which usage is divided into ten-minute blocks from the 5 PM rate drop until 8 PM. All agents prefer to use the phone network from 5:00 to 5:10, having strictly monotonically-decreasing valuations for later slots as compared to earlier slots. Given that time slots are priced identically, rational agents would all choose to use the network from 5:00 to 5:10, leading to a focused load. More formally, consider the operation of a network over t time slots, where each slot has a fixed usage cost of m, and where n risk-neutral agents, $a_1 \dots a_n$, intend to use the network. Agent a_i's valuation for slot s is given by an arbitrary function $v_i(s)$. Let v^l and v^u be lower and upper bounds on every agent's valuation respectively: i.e., $\forall i, s\ v^l(s) \leq v_i(s) \leq v^u(s)$. Thus (for all i) let $\overline{s} = \arg\max_s v_i(s)$ and $\underline{s} = \arg\min_s v_i(s)$. In sections 3 and 4 we make the assumption that $v^l = v^u$ and hence that all agents have the same valuations for all slots (here we use the notation v rather than v_i to describe agents' valuations); we relax this assumption in sections 5 and 6. Agents may also have "names"—numerical identifiers—denoted $name(a_i)$.

To spread out the focused load, the network will provide agents with an incentive to choose other slots. In this paper we consider mechanisms in which agents are spared the usage cost for the slot they choose according to a probability depending on the slot chosen and independent of the probabilities corresponding to other slots. More formally, to prevent all agents from choosing $\overline{s}$, the network implements a mechanism Φ to waive the usage cost for q of the t slots, on average. Free slots will be chosen according to a probability distribution associated with each slot s, which we call $p(s)$. The distribution of agents is denoted d, and so $d(s)$ is the number of agents who chose slot s.

The mechanism implemented by the network specifies p; the network must draw (independently) from each $p(s)$ to determine if the usage cost will actually be waived for slot s.

The number of free slots, q, is thus only an expected value, and there is no guarantee on the total number of slots that will actually be free. Observe that $q = \sum_s p(s)$.

We restrict ourselves to considering mechanisms in which participation is rational for all agents who want to use the network, in all equilibria arising from that mechanism; thus we assume that all agents will participate. An implication of this restriction is that in all equilibria, the expected cost of using the network for any agent must be less than or equal to that agent's valuation for the slot he chooses. Finally, we assume that agents are risk neutral.

In our model, agents must simply choose a slot s. The space of agent strategies may be seen as the space of all functions mapping from the information available to a probability distribution over slot choices. We assume that agents are aware of the mechanism and consider it when determining their strategies. A given agent's expected utility for choosing slot s is $u_i(s) = v_i(s) - (1 - p(s))m$.

It might appear that more powerful mechanisms could be designed if prices could be varied arbitrarily, as opposed to slots priced at either m or 0. In fact, risk-neutral agents are indifferent between any slot priced on the range $[0, m]$ and the same slot made free with an appropriate probability.

2.1 Evaluating Outcomes

The network has two goals: to balance the load caused by the agents' selection of slots and to collect as much revenue as possible. We denote the network's expected revenue given a mechanism Φ and distribution d as $E[R|\Phi, d]$. The network collects a payment of m from each participating agent except for those who receive free slots. As the number of agents is fixed and the mechanism is constrained so that participation is rational for all agents, expected revenue depends only on the likelihood of the usage fee being waived for each slot and on which slots agents select. We define g as the monetary value to the network of the variance of load across the set of time slots. Lower variance corresponds to a more even load and so has a higher dollar value; thus g must decrease strictly as variance increases. We will that load is *balanced* when g is maximized, which corresponds to minimal variance. We define the *superlinear summation* class of functions to be the set of functions in which $g(d) = -\kappa \sum_i h(d(i))$, where h is superlinear in $d(i)$ and κ is a constant that is used to indicate the relative value of load balancing to the network.

Maximizing revenue and maximizing g are conflicting goals, as it costs the network more to induce an agent to choose slot $\underline{s}$ than it does to induce an agent to choose slot $\overline{s}$. (Indeed, note that revenue is maximized when all agents choose $\overline{s}$—i.e., under focused loading—because agents are willing to distribute themselves in this way without the application of any external incentives.) The network must therefore trade off quality of load balancing against expected revenue; the degree of trade-off desired may be specified through the choice of κ. We define z, the evaluation of distribution d under equilibrium φ of mechanism Φ : $z(\Phi, d) = E[R|\Phi, d] + g(d)$. First we define optimality:

Definition 1. *A mechanism-equilibrium pair (Φ, φ) is optimal if and only if for all other pairs (Φ', φ') and for all d, d' resulting from φ and φ' respectively, $z(\Phi, d) \geq z(\Phi', d')$, where n is held constant.*

This definition of optimality is inappropriate for the case where agents have different valuation functions that are not known by the network—the case we take up in sections 5 and 6. In the event that the network's bounds on agents' valuations are not tight, the best mechanism that the network can choose will not extract the maximal amount of revenue from each agent, and so will not be optimal as defined above. Instead, we provide an alternate notion of optimality that bounds the average loss per agent.

Definition 2. *A mechanism-equilibrium pair (Φ, φ) is c-optimal if and only if for all other pairs (Φ', φ') and for all d, d' resulting from φ and φ' respectively, $z(\Phi, d) + cn \geq z(\Phi', d')$, where n is held constant and $c > 0$.*

We also use the term *optimal* to refer to equilibria alone when the mechanism giving rise to the equilibrium is unambiguous. Finally, we give a definition to describe the best possible distribution of agents given a mechanism. A distribution is *ideal* if it maximizes z given the mechanism.

Definition 3. *A distribution d is ideal for mechanism Φ if and only if $d = \arg\max_{d'} z(\Phi, d')$. In such a case we superscript d as d^* to highlight the fact that it is ideal .*

3. PENNY MATCHING

Here we consider a simple mechanism designed to make agents indifferent between all time slots despite their initial preferences. We call it 'penny matching', since agents must guess what slots the network will make free; more formally, we denote this mechanism as Φ_1 . The mechanism follows:

1. Free slots are determined by drawing from p.
2. Agents choose a slot.

All things being equal, agents prefer slot $\overline{s}$ to slot $\underline{s}$. We can overcome this preference by biasing $p(s)$. Recall that an agent's expected utility is given by $u_i(s) = v(s) - (1 - p(s))m$. We can make agents indifferent between slots by requiring that all time slots will have the same expected utility for agents: that is, that the expected utility derived from each time slot is equal to the average expected utility over all time slots. This is expressed by the equation $v(s) - (1 - p(s))m = \frac{1}{t}\sum_i (v(i) - (1 - p(i))m)$. Rearranging, we get:

$$p^*(s) = \frac{\frac{1}{t}(qm + \sum_i v(i)) - v(s)}{m} \tag{1}$$

If free slots are awarded according to p^*, it is a weak equilibrium for all agents to select a slot uniformly at random. We call this equilibrium φ_1 . Consider the case where all other agents play according to φ_1 , and one remaining agent a_i must decide his strategy. Since the choice of any slot entails equal utility on expectation, a_i can do no better than to randomly pick a slot. φ_1 is a weak equilibrium: indeed, there is no strategy that would make a_i worse off.

It appears that deviation from φ_1 will never be profitable for agents, since we have guaranteed that all slots provide the same expected utility. Consider the most profitable deviation, from $\underline{s}$ to $\overline{s}$. We have claimed that utility of both slots is the same: $v(\overline{s}) - (1 - p(\overline{s}))m = v(\underline{s}) - (1 - p(\underline{s}))m$. Since we want to interpret $p(\underline{s})$ and $p(\overline{s})$ as probability measures, $p(\overline{s}) \geq 0$ and $p(\underline{s}) \leq 1$. Substituting the constraints into equation (1), we get $\frac{tv(\overline{s}) - \Sigma_i v(i)}{m} \leq q \leq \frac{t(v(\underline{s}) + m) - \Sigma_i v(i)}{m}$. We must also ensure that a value of q exists for a given m and v. Intersecting the two bounds and simplifying gives us

$m \geq v(\overline{s}) - v(\underline{s})$. We now show how the network can maximize revenue. We define v_{avg} as $\frac{1}{t}\sum_s v(s)$. The requirement that an agent's utility for slot s must be greater than or equal to zero—i.e., that $v(s) - (1 - p(s))m \geq 0$—can be rewritten, substituting in p^*, as $v_{avg} - (1 - \frac{q}{t})m \geq 0$. The seller's revenue will be maximized when all agents get zero utility. Thus we must have $(1 - \frac{q}{t})m = v_{avg}$. There is a range of q and m values that will satisfy this equation; here we show one. We substitute in the lower bound for q from section 3. Rearranging, we get $m = v(\overline{s})$. This satisfies the constraint on m, so we are done. Intuitively, we have shown that we can collect maximum revenue: we can ensure that on expectation each agent will pay an amount exactly equal to his utility for any slot he chooses. However, φ_1 is not guaranteed to achieve an optimal distribution of agents, and therefore, φ_1 is not optimal. The easiest way to show this is to present another equilibrium of Φ_1 that *is* optimal.

Consider an equilibrium in which each of the agents deterministically chooses one slot. (Recall that *any* strategy is rational under Φ_1 , and thus that any set of strategies is a weak equilibrium.) In one such equilibrium, agents deterministically choose slots so that the distribution of all agents is ideal ; we call this equilibrium φ_1^* . Unsurprisingly:

THEOREM 1. (Φ_1, φ_1^*) *is optimal.*

All proofs are deferred to the full version of the paper.

φ_1^* is optimal, but it is extremely unlikely that this equilibrium would arise through the choices of real agents. This drawback is inherent to the setting as we have modeled it so far; a "matching pennies" mechanism can only yield weak equilibria. In the next three sections, we explore more complex mechanisms that give rise to strict equilibria.

4. BULLETIN BOARD SYSTEM

In this section we assume that agents are given a *bulletin board system*: a forum in which all communications are visible to all agents and the identity of agents is associated with their transmissions. For simplicity, we allow a very limited form of communication: agents sequentially indicate the slot that they intend to choose. Let $d_b(s)$ denote the number of agents who have indicated that they will choose slot s. Agents' communications through the bulletin board are *cheap talk*: a technical term that indicates that these communications are not binding in any way. Even so, the bulletin board can help agents to coordinate on desirable equilibria without the use of names. Mechanism Φ_2 follows:

1. "Potentially free"[1] slots chosen according to $(1+\varepsilon)p^*$.
2. Agents communicate through the bulletin board.
3. Agents choose time slots.
4. If $d = d^*$, then "potentially free" slots are made to be free. Otherwise, all agents are made to pay.

A strict equilibrium in Φ_2 , called φ_2 , is for the i^{th} agent to choose a slot s such that $d_{i-1}(s) < d_i^*(s)$; to indicate his chosen slot s on the bulletin board; and ultimately to choose that slot s. Consider the case where all other agents follow φ_2 and agent a_i must decide his strategy. If a_i cooperates and chooses slot s then the distribution of agents is guaranteed to be d^* and so a_i will receive an expected utility of $v(s) - (1 - (1+\varepsilon)p^*(s))m$. If a_i defects to slot s', one of two cases will result. In the first case, agents indicating their choices after a_i will compensate for his deviation by choosing different slots; thus a_i will receive the same expected utility as he would have received if he had not deviated. In the second case, a_i will be late enough in the sequence of agents indicating their choices that the agents who indicate after him will be too few to bring the distribution back to d^*. In this case a_i will receive an expected utility of $v(s') - m$. Since a_i does not know the total number of agents, he must assign non-zero probability to the second case, regardless of the number of agents who have already indicated. Therefore φ_2 is strict as long as $v(s) + (1+\varepsilon)p^*(s)m > v(s')$ for all s, s' such that $1 \leq s, s' \leq t$. Simplifying, we derive the same conditions on q and m described above, except that the probability of a free slot is increased to make φ_2 strict.

It is well known that any game having an equilibrium arising from cheap talk coordination has other equilibria in which agents ignore the cheap talk [2]. Φ_2 is no exception. All agents choosing $\overline{s}$ (focused loading) is an equilibrium when the resulting d could not be transformed into d^* by one agent choosing a different slot. Note, however, that φ_2 Pareto-dominates all equilibria where the cheap talk is ignored. Because of this equilibrium, we know that m cannot be set above $v(\overline{s})$, because agents would receive negative utility in equilibrium. For this reason, we can do no better than setting m as in section 3. Note that there can never be an equilibrium in which participation is irrational if $m \leq v(\overline{s})$, because agents who choose slot $\overline{s}$ will always have nonnegative utility.

THEOREM 2. *There does not exist an optimal* (Φ, φ) *for which* φ *is a strict equilibrium and* $m \leq v(\overline{s})$.

However, there exists no equilibrium of any other mechanism yielding z larger than $z(\Phi_2, \varphi_2) + \varepsilon$.

THEOREM 3. (Φ_2, φ_2) *is* ε*-optimal.*

Note that in φ_2 each agent chooses a slot that would result in an optimal distribution if he were the last agent to post to the bulletin board. In the full version of the paper, we show that we can assign names to agents greedily with the guarantee of achieving the ideal distribution for whatever number of agents eventually participate.

5. COLLECTIVE REWARD

We now consider the more general case where each agent may have a different v_i, bounded by v^l and v^u, as described in section 2. In this section we introduce the assignment of agent names as a mechanism for the agents to coordinate to a desirable equilibrium, and also show how collective reward may be used to prevent agents from deviating. We define mechanism Φ_3 as follows:

1. Each agent indicates that he will participate.
2. Integral names are assigned to agents from $[1, t]$.
3. Each agent indicates what slot he selects.
4. After all agents have selected their slots, the network determines whether each slot will be made free.

The chance that slot s will be free, $p(s)$, depends on the number of agents who chose that slot, $d(s)$. Let *count*(s) be

[1] We redefine q as the expected number of "potentially free" slots; the same redefinition is required for section 6.

the number of agents who were given the name s. Define $d^+(s) = d(s) - count(s)$. For Φ_3 :

$$p(s) = \begin{cases} p^l(s) & \text{if } d^+(s) \leq 0 \\ 0 & \text{if } d^+(s) > 0 \end{cases} \quad (2)$$

Intuitively, we construct p^l so that each agent a_i will participate in the worst case for Φ_3 : when a_i has the lowest possible valuation for the slot corresponding to his name, and the highest possible valuation for all other slots. The derivation of p^l ensures that no agent will deviate regardless of his actual v. We follow the derivation of p^*, with some changes. The equation to make agents indifferent between all slots is changed to: $v^l(s)-(1-p^l(s))m = \frac{1}{t}\sum_i(v^u(i)-(1-p^l(i))m)$. The left-hand side uses v^l so that it represents the lowest possible value for the expected utility of a slot s that the agent could receive free. The right-hand side uses v^u because it represents the most an agent can receive by choosing another slot. If this equality holds, then for all possible v functions for an agent, he will not have incentive to deviate. Algebraic manipulation gives:

$$p^l(s) = \frac{\frac{1}{t}(qm + \sum_i v^u(i)) - v^l(s)}{m} \quad (3)$$

As in section 3, we can derive bounds on q and m. In this case the most profitable possible deviation is from $\underline{s}$ with a valuation of $v^l(\underline{s})$ to $\overline{s}$ with a valuation of $v^u(\overline{s})$. This leads to the following natural condition on m which shows us how to create a large enough p^l to ensure that no agent deviates: $m \geq v^u(\overline{s}) - v^l(\underline{s})$. It also follows that the inequality $v^u(\overline{s}) - m \geq v^l(\underline{s}) - (1 - p^l(\underline{s}))m$ must hold. Substituting in the additional constraint of $p^l(\overline{s}) \geq 0$ and rearranging, we get $q \geq \frac{tv^l(\overline{s})-\Sigma_i v^u(i)}{m}$. Φ_3 sets $m > v^u(\overline{s})$ so that it is never rational for an agent to deviate. We then plug m into the bound on q given above and set q as small as possible to maximize expected revenue.

An equilibrium φ_3 is for each agent a_j to select the slot corresponding to its number. Consider the case where all other agents follow this strategy, and one remaining agent a_i decides his strategy. If agent a_i selects slot s as above then his expected utility is $u_i(s) = v_i(s)-(1-p^l(s))m$. Deviating to slot s' gives him $u_i(s') = v_i(s') - m$. The difference between these two options is $u_i(s) - u_i(s')$, which simplifies to at least the sum of two positive terms: $(v_i(s) - v^l(s)) + (v^u(\overline{s}) - v_i(s'))$. Since agents can only lose by deviating, φ_3 is a strict equilibrium.

There are no equilibria of Φ_3 for which $d \neq d^*$. Consider any distribution of agents such that $d \neq d^*$. There must be some s_1 such that $d^+(s_1) < 0$, and some other s_2 such that $d^+(s_2) > 0$. An agent in s_2 thus has no chance of a free slot, and he receives negative utility for this slot, because $m > v^u(\overline{s})$. If he switches to s_1, then his probability of receiving a free slot becomes $p^l(s_1)$ because $d^+(s_1) \leq 0$. Since p^l is constructed to make participation rational, the agent must receive nonnegative expected utility for this slot, contradicting the claim that staying in s_2 was an equilibrium.

THEOREM 4. *φ_3 is c-optimal,* $c = \max_s(v^u(s) - v^l(s))$.

6. DISCRIMINATORY MECHANISM

A disadvantage of both Φ_2 and Φ_3 is that they reimburse some agents at the end of the game rather than simply waiving their fees. This requires tracking individual agents' behavior and executing more financial transactions, both of which could be costly to the network. Also, Φ_2 has non-optimal equilibria. Finally, irrational agents can harm others in both Φ_2 and Φ_3 . These problems are eliminated by Φ_4 , which makes use of agent names and also discriminates by offering different free slots to different agents:

1. Each agent indicates that he will participate.
2. Integral names are assigned to agents from $[1, t]$.
3. "Potentially free" slots are chosen according to p^l.
4. Each agent indicates what slot he selects.
5. The network checks only those agents in each slot s_i that was picked to be "potentially free". If agent a_j in slot s_i has $name(a_j) = s_i$ then he receives the free slot; otherwise he is made to pay.

Agent a_i's dominant strategy is to choose the slot that may be free for him. The analysis is the same as for φ_3 ; we call this equilibrium φ_4 . Since φ_4 results from (strongly) dominant strategies, it is unique. By theorem (4), d^l is c-optimal for Φ_4 , $c = \max_s(v^u(s) - v^l(s))$.

It may seem disappointing from a game-theoretic point of view that neither strategy nor even payoffs under Φ_4 depend on the actions of other agents. However, we feel that this is an advantage of Φ_4 , since irrational agents are not able to harm others, retroactive payments to agents are not required, and the only equilibrium that exists is c-optimal.

7. CONCLUSION

Focused loading is a predictable problem that occurs in real networks. We present a theoretical model of the problem and discuss four mechanisms that induce selfish agents to smooth out their resource demands. We show a simple mechanism that achieves a weak load-balancing equilibrium, and three more complex mechanisms that balance load in strict equilibria or dominant strategies. Two of our mechanisms assume that all agents value time slots identically, and two generalize to the case where only upper and lower bounds are known on agent valuations. In the future we intend to examine the cases where agents have unrestricted valuations for time slots, and make resource demands of different magnitudes.

8. REFERENCES

[1] D. Blair. Impact of the Internet on core switching network. In *Proc. of ENPW'98*, Les Arcs, France, March 1998.

[2] V.P. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, November 1982.

[3] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queuing algorithm. *Journal of internetworking research and experience*, pages 3–26, October 1990.

[4] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.

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

[6] J.K. MacKie-Mason and H. Varian. Pricing the internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*. Prentice-Hall, 1994.

[7] B. Mitchell. Pricing policies in selected European telephone systems. In *Proceedings of 6th Conference on Telecommunications Policy Research*, pages 437–475, 1978.

[8] A. Odlyzko. A modest proposal for preventing internet congestion. Technical Report TR 97.35.1, AT&T Research, September 1997.

[9] S. Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking*, 3:819–831, 1995.