6.454 : Network Bandwidth Allocation and Stability

Anthony Accardi

October 3, 2006

1 Purpose

We summarize some recent published research on the problem of how to allocate bandwidth to users in a network. While all the references cited at the end were considered, we draw most heavily from the work in [1] and [7].

2 Network Terminology

Let's being with a physical situation to motivate the problem.

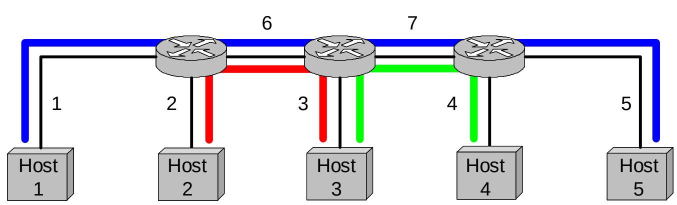

Figure 1: A network diagram, showing links (black) and routes (color).

Say we have hosts and routers arranged in a network, connected with 7 *links*, each of which has a bandwidth capacity C_j. We consider 3 paths through the network, called *routes*. One or more *users* can move data along each route, but the rate cannot exceed the route's allocated bandwidth Λ_i. We wish to determine the optimal bandwidth allocation to maximize some notion of performance. We will model this network, and then consider various performance criteria.

We start by throwing away parts not relevant to the problem, such as the hosts and routers. The *links*, which we call more generally *resources*, are of prime importance. A route becomes nothing more than a set of resources:

how the links were connected is not important, and only the fact that a route must consume the same bitrate on each link is relevant. We can capture this information using a matrix A, whose $(j,i)^{\text{th}}$ entry indicates whether resource j is used by route i.

Each route is used by a number of *users* (aka *flows*) N_i, which varies with time. We assume that the allocated bandwidth Λ_i for a route is divided equally among its users. We are left with the equivalent representations shown in Figure 2.

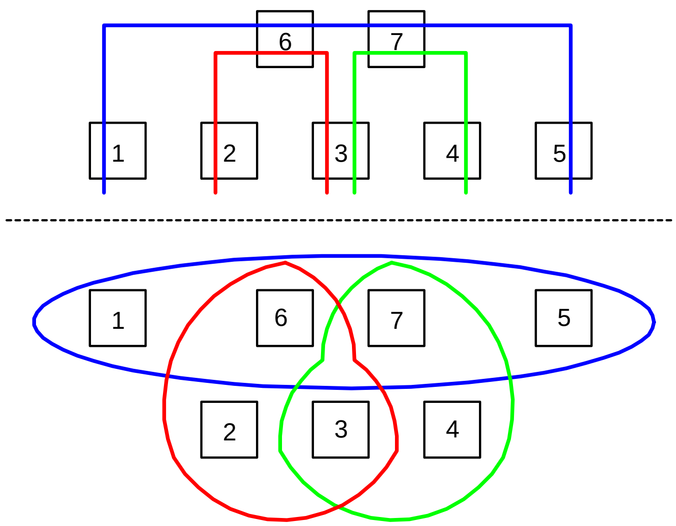

Figure 2: Two abstract views of resource usage.

We summarize the quantities introduced thus far.

C_j	:	capacity of resource (link) j
$N_i(t)$	:	number of users (flows) on route i at time t
$\Lambda_i(\mathbf{N}(t))$	:	allocated bandwidth for route i, depending on network usage at time t
$A = [A_{ji}]$	:	$A_{ji} = 1$ if resource j is used by route i, and $A_{ji} = 0$ otherwise

Using these, we define the capacity constraint

$$\sum_{i\,:\,N_i(t)>0} A_{ji}\Lambda_i(\mathbf{N}(t)) \leq C_j,$$

which must hold for all resources j and for all time t. A bandwidth allocation $\mathbf{\Lambda}$ that satisfies the capacity constraints is called feasible.

3 A Stochastic Model

We wish to study dynamic network behavior, so we model each $N_i(t)$ as a stochastic process. New users arrive according to a Poisson process of rate ν_i, and begin transferring a document with size drawn from an exponential distribution with parameter μ_i. We define the traffic intensity on route i as $\rho_i = \nu_i/\mu_i$. Now $\mathbf{N}(t)$ is a Markov process with a countable state space, a small portion of which is shown in Figure 3 for an example with 3 routes.

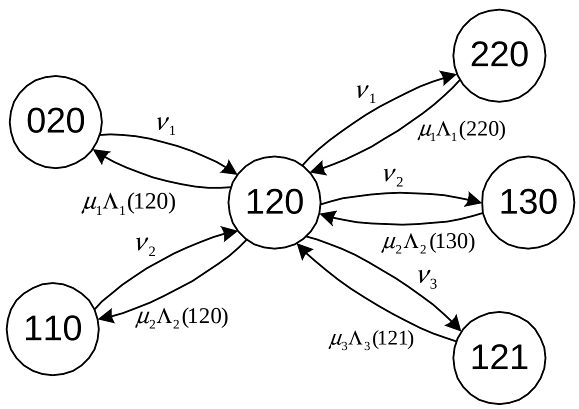

Figure 3: Markov state transition example. State 120 indicates that there are 1, 2, and 0 users on routes 1, 2, and 3 respectively.

4 Performance Criteria

There are many desirable properties and behavior to choose from, and picking the right ones to focus on is a challenge.

4.1 Efficiency

An efficient allocation is one where no bandwidth is wasted. A feasible $\mathbf{\Lambda}$ is efficient if we don't have $\boldsymbol{\mu} \geq \mathbf{\Lambda}$ for any other feasible allocation $\boldsymbol{\mu}$. That is, in an efficient allocation, there is a link along every route that is at capacity.

While efficient allocations are desirable, there are many such allocations for a given network. Efficiency will serve as a sanity check for more complicated schemes as we search for more discriminating criteria.

4.2 Stability, Utilization, and Throughput

Stability is a deeper property. We want the number of flows in progress on each route to remain finite with probability 1. A necessary condition is that

$$\sum_i A_{ji}\rho_i < C_j \tag{1}$$

for all resources j, since if one of these constraints is violated, the corresponding resource's average rate of completing active flows would not exceed the average rate of new flows arriving. In fact, this inequality says something about the utilization of a resource. While necessary, this condition is not sufficient for stability. We will therefore have to assess stability on a case-by-case basis.

Maximizing overall user throughput $\sum_{i:N_i(t)>0} \Lambda_i$ may seem like a good idea. But, there need not be a unique allocation that achieves this. Furthermore, some allocations can produce unexpected results. Consider the simple linear network in Figure 4, where each resource has unit capacity $C_j = 1$.

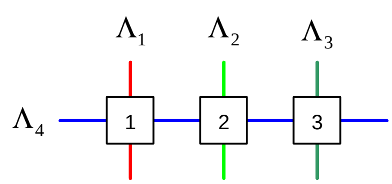

Figure 4: A linear network with 3 resources and 4 routes.

One way to maximize throughput at every point in time is to always allow the short routes to preempt the long route. Assuming that the flow arrivals at different routes are independent, we see that we must add another condition to equation (1) for stability: $\rho_4 < \prod_{i=1}^{3}(1-\rho_i)$.

This is because the long route can only sustain a non-zero rate of 1 for a fraction $\prod_{i=1}^{3}(1-\rho_i)$ of the time. Now, if we set $\rho_i = 1/3$, we see that equation (1) is easily satisfied, and in fact no resources are near fully utilized, yet the network is not stable. Similar behavior is exhibited by other class-based queueing schemes in more complicated networks. There is a utilization price to pay for schemes that give preferential treatment to one route over another.

4.3 Max-Min Fairness

We now direct our attention to fair allocation schemes. We increase the allocation for each user, unless doing so requires a corresponding decrease for a user of equal or lower bandwidth to satisfy the capacity constraints. The resulting allocation is uniquely determined and called max-min fair. There is a simple greedy algorithm for computing it: starting with an empty network, increase the allocation for all flows equally until a resource constraint is hit. Fix the allocations for routes passing through this bottleneck, and repeat the procedure with the rest.

Max-min fair allocations prioritize fairness over utilization and require global information to implement. Specifically, consider the example in Figure 5, where a max-min allocation gives somewhat undesirable results.

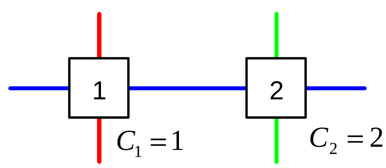

Figure 5: Example where max-min fairness sacrifices utilization.

4.4 Proportional Fairness

Proportional fairness seeks to remedy this by considering the proportional increase in bandwidth for each route. Specifically, a feasible allocation $\mathbf{\Lambda}$ is proportionally fair if for any other feasible allocation $\mathbf{\Lambda}^*$ we have

$$\sum_i N_i(t) \frac{\Lambda_i^* - \Lambda_i}{\Lambda_i} \leq 0.$$

This corresponds to maximizing a logarithmic utility function $\sum_i N_i(t) \log \Lambda_i$, since a small perturbation $\mathbf{\Lambda} + \mathbf{\Delta}$ would increase the utility function by $\sum_i N_i(t) \frac{\Delta_i}{\Lambda_i}$. Moreover, distributed algorithms to implement proportional fairness are known [6].

4.5 α-Fair Allocations

The utility maximization interpretation of proportional fairness motivates generalizing to weighted α-fair allocations, which are solutions to the maximization of

$$G(\mathbf{\Lambda}, t) = \begin{cases} \sum_{i\,:\,N_i(t)>0} \kappa_i N_i^{\alpha}(t) \frac{\Lambda_i^{1-\alpha}}{1-\alpha} & \text{if } \alpha \in (0,\infty),\ \alpha \neq 1 \\ \sum_{i\,:\,N_i(t)>0} \kappa_i N_i \log \Lambda_i & \text{if } \alpha = 1 \end{cases}, \qquad (2)$$

subject to the capacity constraints

$$\sum_{i\,:\,N_i(t)>0} A_{ji}\Lambda_i \leq C_j,$$

over

$$\Lambda_i \geq 0.$$

Various special cases are evident for specific values of α and the weights κ_i. First set $\kappa_i = 1$. When $\alpha \to 0$, we are just maximizing throughput. When $\alpha \to \infty$, the smallest allocations dominate the objective function G. This leads to max-min fairness. Setting $\alpha = 1$ gives proportional fairness, seemingly by definition of the objective function. But this is just a technicality – maximizing the bottom part of equation (2) when $\alpha = 1$ is equivalent to maximizing the top for $\alpha \approx 1$. To see this, expand about $\Lambda_i = 1$ using a Taylor series:

$$\frac{\Lambda_i^{1-\alpha}}{1-\alpha} = \frac{1}{1-\alpha} + \frac{1}{1!}(\Lambda_i - 1) - \frac{\alpha}{2!}(\Lambda_i - 1)^2 + \frac{\alpha(1+\alpha)}{3!}(\Lambda_i - 1)^3 + \cdots.$$

The leading constant $\frac{1}{1-\alpha}$ is irrelevant to the optimization, so we drop it. Setting $\alpha = 1$ then gives us $\log \Lambda_i$.

TCP corresponds to $\alpha = 2$ and $\kappa_i = \frac{1}{\mathrm{RTT}_i^2}$, where RTT_i is the round-trip transit time on route i. This is because TCP is biased against long round-trip delays. A TCP source sets the size of a window of data which it will attempt to send ahead of what the receiver has acknowledged. This window is incremented in size (up to a maximum) whenever an acknowledgment is received, which happens once every RTT.

4.6 α-Fair Allocation Properties

α-fair allocations have a number of useful properties. We assume that $N_i(t) > 0$ for simplicity.

1. Uniqueness and positivity. The constraints define a compact convex set. Treat the cases $\alpha \in (0, 1)$ and $\alpha \in [1, \infty)$ separately. In both cases, the objective function G is continuous and strictly concave in $\mathbf{\Lambda}$, so there is a unique maximum. In the first case $\frac{\partial G}{\partial \Lambda_i} \to +\infty$ as $\Lambda_i \to 0$, and in the second $G \to -\infty$ when $\Lambda_i \to 0$, so an optimum solution must have strictly positive components.

2. Scaling. The solution is invariant to scaling: $\mathbf{\Lambda}(r\mathbf{N}) = \mathbf{\Lambda}(\mathbf{N})$ for all $r > 0$. Scaling $\mathbf{N}$ by r amounts to scaling G by r^α, a constant. The optimal allocation does not change.

3. Continuity. The optimal solution for $\mathbf{\Lambda}$ is continuous in $\mathbf{N}$. A proof requires some work, but essentially follows by continuity of the other quantities and uniqueness.

The big question is what can we say about the stability of an α-fair allocation? It turns out that equation (1) becomes a necessary and sufficient condition for stability in this case. The result even holds under more general modeling conditions, without assumptions on Poisson arrivals and exponential document sizes [4].

5 A Fluid Model to Prove Stability

In theory, given joint distributions for the $N_i(t)$, we can solve the optimization problem for the distribution of the optimal α-fair allocation $\mathbf{\Lambda}(t)$, and then answer the stability question. But this is very difficult to do, unless the distribution for $N_i(t)$ is very simple (e.g., Poisson). But simplifying our model in this way ignores one of the key features we are attempting to capture: how service time is impacted under high load.

Our strategy is therefore to simplify our stochastic model in another way, by approximating it with a fluid model. The result is easier to work with, yet still captures the behavior we are interested in. We can prove results such as stability with the fluid model, and then carry them over to our original model.

A fluid model solution is an absolutely continuous function $\bar{\boldsymbol{N}}(t)$ so that at each regular point t and each route i

$$\frac{\mathrm{d}}{\mathrm{d}t}\bar{N}_i(t) = \begin{cases} \nu_i - \mu_i \Lambda_i(\bar{\boldsymbol{N}}(t)) & \text{if } \bar{N}_i(t) > 0 \\ 0 & \text{if } \bar{N}_i(t) = 0 \end{cases}, \tag{3}$$

and for each resource j

$$\sum_{i:\bar{N}_i(t)>0} A_{ji}\Lambda_i(\bar{\boldsymbol{N}}(t)) + \sum_{i:\bar{N}_i(t)=0} A_{ji}\rho_i \leq C_j. \tag{4}$$

Equation (3) captures the infinitesimal drift of our original stochastic processes $N_i(t)$. We can interpret equation (4) as saying that the allocated bandwidth cannot exceed the capacity for any resource, whereas the allocation of any currently unused route is considered to be ρ_i.

To connect the two models, we decompose our Markov process into non-decreasing counting processes:

$$N_i(t) = N_i(0) + E_i(t) - S_i(T_i(t)),$$

which track the initial state, arrivals, and departures separately. For example, $E_i(t)$ is a Poisson process with rate ν_i. We scale each process, obtaining a family $\frac{E_i(rt)}{r}$ indexed by r. As $r \to \infty$, we can apply tools such as the strong law of large numbers for renewal processes [3] to conclude, for example, that $\frac{E_i(rt)}{r} \to \nu_i t$ with probability 1. Together with the scaling and continuity properties discussed above, one can show how a sequence of scaled stochastic models converges to a fluid model solution. (See Appendix B of [7] for further details.)

Now, one can show that if a fluid model solution empties in finite time, $\bar{\boldsymbol{N}}(t) = 0$ for $t \geq T$, then the original Markov process is positive recurrent [2]. Finally, a Lyapunov function can be found that shows that this indeed happens [1].

6 Further Reading

The 4 primary papers are [1], [7], [8], and [4]. Bonald and Massoulié give a concise overview of the problem and main stability result in [1]. However, their argument for generalizing the model is wrong. Kelly and Williams describe the fluid model in detail in [7]. Massoulié uses duality theory and convex analysis to give an alternative proof of stability for proportional fairness in [8]. Gromoll and Williams correctly prove the general stability result using a measure theoretic argument in [4].

There are other, older papers listed below that provide more background. Dai shows how fluid model properties carry over to the original model in [2]. Kelly motivates proportional fairness in [5], and specifies algorithms in [6]. Finally, Mo and Walrand assess TCP fairness in [9].

References

[1] T. Bonald and L. Massoulié, “Impact of Fairness on Internet Performance,” *Proceedings of ACM Sigmetrics*, 2001.

[2] J. G. Dai, “On Positive Harris Recurrence of Multiclass Queueing Networks: a Unified Approach via Fluid Limit Models,” *The Annals of Applied Probability*, Vol. 5, No. 1, 1995, pp. 49-77.

[3] R. G. Gallager, *Discrete Stochastic Processes*, Boston: Kluwer Academic Publishers, 1996.

[4] H. C. Gromoll and R. J. Williams, “Fluid Limit of a Network with Fair Bandwidth Sharing and General Document Size Distributions,” 2006. Available http://math.ucsd.edu/ williams/bandwidth/gendocsize.html.

[5] F. P. Kelly, “Charging and Rate Control for Elastic Traffic,” *European Transactions on Telecommunications*, Vol. 8, 1997, pp. 33-37.

[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*, 1998, pp. 237-252.

[7] F. P. Kelly and R. J. Williams, “Fluid Model for a Network Operating Under a Fair Bandwidth-Sharing Policy,” *The Annals of Applied Probability*, Vol. 14, No. 3, 2004, pp. 1055-1083.

[8] L. Massoulié, “Structural Properties of Proportional Fairness: Stability and Insensitivity,” 2005. Available ftp://ftp.research.microsoft.com/pub/tr/TR-2005-102.pdf.

[9] J. Mo and J. Walrand, “Fair End-to-End Window-Based Congestion Control,” *IEEE Transactions on Networking*, Vol. 8, No. 5, 2000, pp. 556-567.