

# Admission Control to Minimize Rejections

Avrim Blum, Adam Kalai, and Jon Kleinberg

**Abstract.** Admission control (call control) is a well-studied online problem. We are given a fixed graph with edge capacities, and must process a sequence of calls that arrive over time, accepting some and rejecting others in order to stay within capacity limitations of the network. In the standard theoretical formulation, this problem is analyzed as a benefit problem: The goal is to devise an online algorithm that accepts at least a reasonable fraction of the maximum number of calls that could possibly have been accepted in hindsight. This formulation, however, has the property that even algorithms with optimal competitive ratios (typically  $O(\log n)$  where  $n$  is the number of nodes) may end up rejecting the vast majority of calls even when it would have been possible in hindsight to reject only very few.

In this paper, we instead consider the goal of approximately *minimizing* the number of calls *rejected*. This is much more appropriate for settings in which rejections are intended to be rare events. In order to avoid trivial lower-bounds, we assume *preemption* is allowed and that calls are given to the algorithm as fixed paths. We show that in a number of cases, we can in fact achieve a competitive ratio of 2 for rejections (so if the optimal in hindsight rejects 0, then we reject 0; if the optimal rejects  $r$ , then we reject at most  $2r$ ). For other cases, we get worse, but nontrivial, bounds. For the most general case of fixed paths in arbitrary graphs with arbitrary edge capacities, we achieve matching  $\Theta(\sqrt{m})$  upper and lower bounds, where  $m$  is the number of edges. We also show a connection between these problems and online versions of the vertex-cover and set-cover problems (our factor-2 results give 2-approximations to slight generalizations of the vertex cover problem, much as [Alon et al. 99] show hardness results for the benefit version based on the hardness of approximability of independent set).

## I. Introduction

In the well-studied *admission control* (or *call control*) problem, our job is to manage a network  $G$  (a graph with edge capacities) in the presence of online requests for communication (calls). Requests for communication may be accepted or rejected, and the goal of an online algorithm is to accept as many as possible while staying within the edge capacities of the network.

This problem has typically been studied as a *benefit problem*. That is, one compares the number of calls that could have been accepted in hindsight to the number actually accepted by the online algorithm, and tries to minimize this ratio. A number of papers have produced good bounds for this metric, such as the work of Awerbuch, Azar, and Plotkin [Awerbuch et al. 93] for the high-capacity setting, and Awerbuch et al. [Awerbuch et al. 94b] for trees and other specific networks. A serious problem with viewing call-control as a benefit problem, however, is that even with, say, an  $O(\log n)$  competitive ratio that would normally be considered quite good, it is possible that the algorithm may route only a  $1/(\log n)$  fraction of the calls even if a solution routing nearly all of them is possible.<sup>1</sup> For many of the natural applications of admission control, even a modest constant fraction of rejections would be deemed unacceptable performance. Thus, for these types of applications, the benefit formulation appears fundamentally flawed.

In this paper, we depart from the benefit metric and instead set our sights on the goal of *minimizing* the number of calls *rejected*. That is, if OPT (the optimal strategy in hindsight) rejects 0, then we should reject 0. If OPT rejects a small number, then we should reject only a small multiple of that. What we show is that for several natural cases, we can in fact achieve a competitive ratio of 2 for rejections. For other versions, we can achieve worse, but still non-trivial, bounds. Of course, approximately minimizing rejections suffers from the reverse problem that the algorithm may accept no calls even if in hindsight it was possible to accept, say, half of them. However, in many applications, even *optimal* performance in such a case would be unacceptable: If one's network required one to reject a significant fraction of calls, then the correct response would be to upgrade the network. It is these types of settings that motivate our work.

We assume in our results that the online algorithm is allowed *preemption*: At any time we may preempt (reject) requests that had previously been accepted,

---

<sup>1</sup>In fact, prior to Leonardi et al. [Leonardi et al. 98], the situation was even worse. Depending on the types of requests made, many of the randomized algorithms would, with probability  $1 - 1/(\log n)$ , accept no calls at all. That is, the variance of possible benefits was high compared to the expectation.

and simply count it as if the request had been rejected from the start. This is one of the standard models and is necessary to achieve any nontrivial bound on rejections. A second assumption we make is that each request is for a unit demand on a fixed path. That is, the requests can be thought of as a sequence of paths  $p_1, p_2, \dots$ , and the decision made by the online algorithm is just whether to accept or reject each path, and does not involve routing. Again, if routing is part of the algorithm's job, then even in very simple settings, no nontrivial bound is possible for our performance metric.<sup>2</sup> Our results are then as follows:

**Admission control on a line.** When the underlying graph is a line, we can achieve a competitive ratio of 2 for any set of edge capacities. That is, the algorithm will reject at most twice as many as the minimum possible in hindsight.

**Admission control on a general graph.** For general graphs, we can achieve a competitive ratio of 2 if all edge capacities are 1 (the disjoint paths case). This extends to a ratio of  $c + 1$  if all edge capacities are  $\leq c$ . For arbitrary capacities, we give a different algorithm that achieves a competitive ratio of  $O(\sqrt{m})$ , where  $m$  is the number of edges, which we match with an  $\Omega(\sqrt{m})$  lower bound.

An interesting aspect of the rejection measure is that the easiest cases are when capacities are *low*. This is the opposite of the situation for the benefit measure, where low capacities are difficult and higher capacities make the problems easier. A preliminary version of this work appears in [Blum et al. 01].

### 1.1. Related Work

As discussed above, existing work on admission control has primarily focused on the problem of maximizing the number of accepted calls, rather than minimizing the number of rejected calls. (See the surveys by Plotkin [Plotkin 95] and Leonardi [Leonardi 98].) The one exception of which we are aware is the work of Kamath, Palmon, and Plotkin [Kamath et al. 96], who provide performance guarantees in terms of a competitive ratio on rejections. Their setting is quite different from ours, however. Most significantly, they assume the input to be probabilistically generated, not worst-case. They consider calls that are generated according to a Poisson process, with exponentially distributed holding

---

<sup>2</sup>Consider a 4-cycle ABCD with capacity  $c$  on each edge. Imagine that we are given  $c$  calls connecting the diagonally opposite nodes A and C, and then we are given either  $c$  calls connecting A and B, or else  $c$  calls connecting A and D, with equal probability. Every online algorithm rejects  $c/2$  calls in expectation, while it was possible to reject none offline. Similarly, with  $n$  separate 4-cycles, the online algorithm rejects  $nc/2$  calls in expectation, while OPT rejects none.

times, and also assume the maximum bandwidth of a call to be very small relative to the available edge capacity (i.e., large capacities). Moreover, their model does not allow preemption.

Routing on fixed paths, as we consider here, was studied by Alon, Arad, and Azar [Alon et al. 99] and Awerbuch, Azar, and Plotkin [Awerbuch et al. 93] under the traditional measure of maximizing benefit. Of course, in linear networks, and more generally in tree networks, one is necessarily routing on fixed paths; for trees, the work of Awerbuch et al. [Awerbuch et al. 94a] provides  $O(\log n)$ -competitive algorithms for maximizing benefit. (See also the earlier work of [Garay and Gopal 92, Garay et al. 93] for linear networks, and the improved probabilistic guarantees obtained by Leonardi et al. [Leonardi et al. 98].)

A number of previous papers have considered the performance gains obtainable by allowing preemption, in the context of maximizing benefit. Adler and Azar [Adler and Azar 99] show that allowing preemption leads to an  $O(1)$ -competitive algorithm for benefit maximization in linear networks.

## 1.2. Notation and Definitions

We are given a graph  $G$ , which may be directed or undirected, with  $m$  edges and  $n$  nodes. Each edge  $e$  has an integer capacity  $c_e > 0$ . We are also given a sequence of requests,  $p_1, p_2, \dots$ , each of which is a simple path in the graph. Each path may either be accepted or rejected. The requirement is that for every edge  $e$ , the number of unrejected requests that have edge  $e$  should be no larger than  $c_e$ . We will call a set of rejection decisions *valid* if it satisfies this requirement.

In the offline problem, we must simply find a small valid set of rejections. In the online problem, we are given requests one at a time, and we must choose to accept or reject the requests online so that the set of accepted requests never exceeds the capacity of any edge. We also allow our online algorithm to preempt an earlier request, i.e., we may reject a request after already accepting it. However, we may not accept a request after rejecting it.

Let  $\text{OPT}$  be a minimum valid set of rejections. We say that an algorithm is  $k$ -competitive ( $k$  may be a function of  $m$ ,  $n$ , and  $c$ ) if the number of requests *rejected* by this algorithm is at most  $k|\text{OPT}|$ .

One final note: Our algorithms will sometimes decide to reject some requests even when not strictly necessary. Because we have preemption, these can always be implemented in a lazy manner. That is, such requests can be marked, but not actually rejected, until a new request arrives that causes a conflict with it.

## 2. Preliminaries: Set-Cover and Vertex-Cover

Imagine a situation where accepting all requests would put every edge one unit over capacity. The offline problem is then exactly set-cover, where each path is represented by the set of edges it contains, and we must choose a set of paths (to reject) that covers each edge at least once. In this case, the online admission control problem can be thought of as a version of online set-cover. While this is not a direct reduction, we hope the following discussion will serve as motivation and intuition for our algorithms later.

A well-known result for the set-cover problem is that if every point is in at most  $k$  sets, then there is a simple  $k$ -approximation algorithm: Pick an arbitrary uncovered point, take all sets that cover it, and repeat. The case  $k = 2$  corresponds to vertex cover.

A slight generalization of the  $k = 2$  case is a setting in which a point may potentially be covered by many sets  $s_1, s_2, \dots$ , but where we are guaranteed that some two of those sets  $s_i, s_j$  cover their union. Then one can achieve a 2-approximation as follows: Pick an arbitrary uncovered point  $p$ , find two sets that cover the union of all sets covering  $p$ , take those two sets, and repeat. This is a 2-approximation because each time two sets are chosen, they can be charged to whatever set  $s_p$  in the optimal solution is used to cover  $p$ . Because the two sets chosen by the algorithm contain  $s_p$ , we are guaranteed that each selection of two sets is charged to a unique set in the optimal cover.

Some of the results below can be viewed as an online version of this algorithm and guarantee.

## 3. Admission Control on a Line

We begin with the special case of a line graph. Each edge  $e$  has some arbitrary capacity  $c_e$ . A request corresponds to an interval on this line and the capacities limit the number of intervals covering any given edge that may be accepted. We show a 2-competitive algorithm, based on the set-cover idea above. The idea is that whenever a new request cannot be accepted due to capacity constraints, we arbitrarily select one of the edges that would go over capacity, and throw out the two requests  $p_l$  and  $p_r$  covering that edge that extend farthest to the left and farthest to the right, respectively. (One of these may or may not be the current request.) We then accept the current request if we did not throw it out. To be more precise:

1. If a request can be accepted, accept it.
2. If a request cannot be accepted, then choose an arbitrary edge  $e$  that would be put over capacity.

- (a) Among the unrejected requests that contain  $e$  (including the current request), let  $p_l$  be one that extends furthest to the left.
  - (b) Among the unrejected requests that contain  $e$  (including the current request), let  $p_r$  be one that extends furthest to the right.
3. Reject  $p_l$  and  $p_r$ , and accept the current request if it is not one of  $\{p_l, p_r\}$ .

**Theorem 3.1.** *The above algorithm is 2-competitive.*

**Proof.** Consider some optimal valid rejection set  $\text{OPT}$ . Each time the algorithm rejects a pair of requests  $\{p_l, p_r\}$ , we will modify  $\text{OPT}$  (call the new set  $\text{OPT}'$ ) by adding at most 1 request to it, in order to maintain an invariant that  $\text{OPT}'$  is a superset of the requests rejected by the online algorithm. We do this as follows. Each time the online algorithm reaches case 2, we know that  $\text{OPT}$  must have rejected at least one request  $p_{opt}$  of those being considered by the online algorithm (i.e., at least one of those covering edge  $e$  that have not yet been rejected by the online algorithm). Therefore, when the online algorithm rejects  $p_l$  and  $p_r$ , we know that (viewing paths as sets of edges)  $p_l \cup p_r \supseteq p_{opt}$ . Therefore, if we put  $p_l$  and  $p_r$  into  $\text{OPT}'$ , and then remove  $p_{opt}$  if neither  $p_l$  nor  $p_r$  had been in  $\text{OPT}'$  already, this only adds 1 to the size of  $\text{OPT}'$ , maintains its status as a valid rejection set, and maintains our invariant. So, if  $t$  is the number of requests rejected by the online algorithm, then  $t \leq |\text{OPT}'| \leq |\text{OPT}| + t/2$ , and therefore,  $t \leq 2|\text{OPT}|$ .  $\square$

Another way of viewing this argument is that each time the algorithm rejects two requests  $p_l$  and  $p_r$ , we give  $\text{OPT}$  a “two-fer,” allowing it to reject those two requests for the price of 1. Since  $\text{OPT}$  must reject some request contained in their union, it might as well take the offer. Inductively, at the end of the game,  $\text{OPT}$  has rejected the exact same set as the online algorithm, but at half the cost.

The above algorithm and analysis also applies if the underlying graph is a cycle (with fixed paths).

## 4. General Graphs

### 4.1. The Low Capacity Case

**Theorem 4.1.** *On a general graph  $G$ , if every edge  $e$  has capacity  $c_e \leq c$ , then there is a simple  $(c + 1)$ -competitive algorithm.*

**Proof.** The algorithm is just an online version of the  $k$ -approximation to set-cover:

1. If a request can be accepted, accept it.
2. If a request cannot be accepted, then choose some edge  $e$  that would be over capacity. Reject the current request along with the  $c_e$  (unrejected) other requests that contain  $e$ .

This algorithm rejects sets of requests of size  $\leq c + 1$  that all share an edge. These sets are disjoint. Any valid rejection set must include at least one request from each of these sets. Therefore, the algorithm achieves a competitive ratio of  $c + 1$ .  $\square$

Thus, if all edges have capacity 1 (the disjoint paths case), we have a 2-competitive algorithm.

## 4.2. General Capacities

The above algorithm gets worse as the capacities in the graph become large. Can we achieve a bound independent of the capacities for general graphs? The connection to set-cover suggests that perhaps we could achieve an  $O(\log m)$  bound. However, it turns out that the online nature of the problem makes that impossible. What we show instead are a set of matching  $\Theta(\sqrt{m})$  upper and lower bounds. We begin with the lower bound.

**Theorem 4.2.** *There is a  $\Omega(\sqrt{m})$  lower bound on the competitive ratio of any online algorithm for general graphs with arbitrary capacities. This holds for randomized algorithms as well.*

**Proof.** For clarity, we will use a multigraph for our lower bound, and then at the end, convert the multigraph into a standard (nonmulti) graph. The multigraph consists of  $k + 1$  vertices  $\{0, 1, \dots, k\}$  arranged in a line, with  $k$  edges connecting each vertex to the next. So the total number of edges is  $k^2$ . Each edge has capacity  $k^{k-1}$ .

We begin by seeing  $k^k$  paths of length  $k$ , one for each possible route between vertex 0 and vertex  $k$ . By design, these will fill all edges exactly to capacity. We then see  $k$  single-edge paths: The first path is a random edge between vertex 0 and vertex 1; the second is a random edge between vertex 1 and vertex 2, and so on.

The offline algorithm needs only to reject one path, namely the path among the first  $k^k$  that happens to match the sequence of  $k$  single-edge paths seen at the end. However, any online algorithm must reject at least  $k/2$  in expectation.

That is because if  $j < k$  paths have been rejected so far, then the next single-edge path seen has at least a  $(k - j)/k$  chance of causing its edge to go over capacity. Hence, as long as  $j < k/2$ , the next edge causes a rejection with probability at least  $1/2$ . Therefore, the competitive ratio of any online algorithm is at least  $k/2$  which is  $\Omega(\sqrt{m})$ .

We can convert the multigraph in the above argument into a standard (non-multi) graph by just putting a node in the middle of each edge, which only doubles the total number of edges.  $\square$

**Remark 4.3.** Our lower bound requires an exponential number of requests and exponential capacity sizes. Perhaps one might be able to give guarantees logarithmic in  $m$  if we also allow logarithmic dependence on the maximum capacity.

We now give a matching  $O(\sqrt{m})$  upper bound. Specifically, we present a  $4\sqrt{m}$ -competitive algorithm for an arbitrary multigraph with  $m$  edges. The algorithm is as follows, starting with zero “chips” on every edge and  $R = 0$ .

1. If a request covers at least  $\sqrt{m}$  edges that have chips, then
  - (a) reject the request.
  - (b) remove one chip from each of the request’s edges (that have chips).
2. If a request cannot be accepted (some edge would be over capacity), then
  - (a) reject the request.
  - (b)  $R = R + 1$ .
  - (c) if  $R$  is a multiple of  $\sqrt{m}$ , then
    - i. add a chip to each edge of the graph.
    - ii. Reapply Step 1 to every accepted request so far.
3. Else, accept the request.

**Theorem 4.4.** *The above algorithm is  $O(\sqrt{m})$ -competitive.*

**4.2.1. Analysis.** Observe that the number of rejections from Step 1 (Line 1(a)), including rejections caused by 2(c)ii) is no more than the number of rejections from Step 2 (Line 2(a)). This is because, after  $R$  Step 2 rejections, we have placed no more than  $R/\sqrt{m}$  chips on each edge, and every Step 1 rejection removes at least  $\sqrt{m}$  chips. Thus, to show that the algorithm is  $4\sqrt{m}$ -competitive, we will show that  $R \leq 2|OPT|\sqrt{m}$ . From here on, when we refer to a rejection, we mean a Step 2 rejection.



Just before we perform each rejection, we “blame” it on an individual edge in one of OPT’s rejections, as follows. Let  $e$  be the first of the edges that would have gone over capacity had we accepted the request. We blame the rejection on the first OPT rejection that we have not yet rejected, has  $e$ , and has not yet been blamed for  $e$ . Not only must there be some such OPT rejection to blame, but it must come no later in the request sequence than the current request. To see this, say we have had  $e$  as a blame edge  $t$  times before, and we have rejected  $r$  of OPT’s rejections that have  $e$ . Then, including the current request, we must have seen  $c_e + r + t + 1$  requests that contain  $e$ . OPT must also reject at least  $r + t + 1$  requests with  $e$ , and we have rejected  $r$  of these and blamed  $t$  of them for  $e$ , leaving at least 1 previous OPT rejection to blame.

Also notice that after  $|OPT|$  chips have been removed from an edge, we will not blame any more rejections on that edge. This is because the total number of requests that have an edge does not exceed  $c_e + |OPT|$ . Finally, it suffices to show that no OPT rejection is blamed for more than  $\sqrt{m}$  rejections after  $R = |OPT|\sqrt{m}$ . At this point, we have placed  $|OPT|$  chips on each edge, where some chips may have been removed. Fix an OPT rejection. Look at the first rejection after  $R = |OPT|\sqrt{m}$  blamed on it. Since we did not reject the OPT rejection in Step 1, it has less than  $\sqrt{m}$  edges with chips. All of its other edges must have had at least  $|OPT|$  chips removed, so they will never be blamed again. Thus, we will blame at most  $\sqrt{m}$  edges on each OPT rejection after  $R = |OPT|\sqrt{m}$ , which implies  $R \leq 2|OPT|\sqrt{m}$  and the total number of rejections is at most  $4|OPT|\sqrt{m}$ .

## 5. The Offline Case

It is interesting to consider the *offline* version of our problem because of its connection to set-cover. For the offline problem, we know the *excess* of each edge, i.e., the number of calls that include that edge minus its capacity. Let  $n_e$  be the excess of edge  $e$ . Our goal is to reject the fewest requests such that each edge  $e$  is contained in at least  $n_e$  rejections. This can be thought of as generalization of a set-cover problem, where each point  $e$  has an associated number  $n_e$ , and instead of the usual goal of covering each point at least once, a legal cover must cover point  $e$  at least  $n_e$  times.

Let us define  $N = \sum_e n_e$ ; that is,  $N$  is the total sum of the excesses. Then, the usual analysis of greedy set-cover gives us an  $O(\log N)$  approximation to this generalized problem. In particular, if we imagine placing  $n_e$  chips on point  $e$ , then the greedy set-cover algorithm becomes: Take the set that covers the most points of those with chips on them, remove one chip from each point covered,

and repeat. If the optimal solution uses  $k$  sets, then at each step, the greedy algorithm must remove at least a  $1/k$  fraction of the chips remaining, giving us the  $O(\log N)$  ratio.

A natural question is whether this upper bound can be improved to  $O(\log m)$  where  $m$  is the number of points (edges). This would be strictly better than what is achievable for the online problem. It is not clear if the greedy algorithm can be used to achieve this, but we *can* get  $O(\log m)$  via randomized rounding as follows.

Formulate the problem as a linear program, where  $0 \leq f_i \leq 1$  is the fraction of set  $s_i$  to take (the fraction of the  $i$ th request to reject). Our objective is to minimize  $\sum f_i$  subject to the chip (capacity) constraints  $\sum_{i:e \in s_i} f_i \geq n_e$ , for every point  $e$ . The minimum value of the objective will be no larger than  $k$ , the value of the optimal integer solution. We solve the LP efficiently and then round the solution. To round, we choose each set independently, with probability of choosing set  $s_i$  equal to  $\min(1, 5f_i \log m)$ .

Next, Chernoff bounds imply that a given point  $e$  is covered at least  $n_e$  times with probability  $\geq 1 - 1/m^2$ . To see this, first note that we do not have to worry about those sets that have  $5f_i \log m \geq 1$ , because we will select them for certain. So ignore these sets. Let  $z_e$  be the sum of the  $f_i$  for the remaining sets  $s_i$  that have  $e$ , and we can assume  $z_e \geq 1$  else we are already done. We expect to select  $5z_e \log m$  sets covering  $e$ , and the only way in which we could fail is if the number we actually select is less than  $z_e$ . By the multiplicative Chernoff bounds,<sup>3</sup> this happens with probability less than

$$e^{-\left(1 - \frac{1}{5 \log m}\right)^2 (5z_e \log m)/2} \leq \frac{1}{m^2}$$

for sufficiently large  $m$ . Thus, with probability at least  $1 - m/m^2$ , we will have *all* points covered the desired number of times. Furthermore, the expected number of sets chosen is  $O(k \log m)$ , as desired.

## 6. Conclusions

We have shown that in a number of natural cases, we can achieve good, or at least nontrivial, bounds for minimizing the number of rejections in admission control. One open question left by these results is that our  $\Omega(\sqrt{m})$  lower bound requires an exponential number of requests and exponential size capacities. Perhaps one might be able to achieve bounds that are logarithmic in  $m$  if we also allow

<sup>3</sup>If  $X$  is a sum of independent  $\{0, 1\}$ -valued random variables, with  $\mu = E[X]$ , the multiplicative Chernoff bound states that  $Pr[X < (1 - \epsilon)\mu] \leq e^{-\epsilon^2 \mu/2}$ .

logarithmic dependence on the maximum capacity  $c$ . Another interesting open problem is whether these results can be extended to the case where different requests may have different demands.

## Acknowledgments

We would like to thank the referees for a number of helpful comments. This research was conducted while the second author was at Carnegie Mellon University, and was supported in part by NSF grants CCR-9732705 and CCR-0085982, NSF Faculty Early Career Development Award CCR-9701399, a David and Lucile Packard Foundation Fellowship, an ONR Young Investigator Award, and an IBM Graduate Fellowship.

## References

- [Adler and Azar 99] R. Adler and Y. Azar. “Beating the Logarithmic Lower Bound: Randomized Preemptive Disjoint Paths and Call Control Algorithms.” In *Proceedings of 10th SIAM-ACM Symposium on Discrete Algorithms*, pp. 1–10, Philadelphia, PA: SIAM, 1999.
- [Alon et al. 99] N. Alon, U. Arad, and Y. Azar. “Independent Sets in Hypergraphs with Applications to Routing via Fixed Paths.” In *Proceedings of 2nd APPROX*, Lecture Notes in Computer Science 1671, pp. 16–27, Heidelberg: Springer-Verlag, 1999.
- [Awerbuch et al. 93] Baruch Awerbuch, Yossi Azar, and Serge Plotkin. “Throughput-Competitive Online Routing.” In *Proc. 34th Symp. Foundations of Computer Science*, pp. 32–40, Los Alamitos, CA: IEEE Press, 1993.
- [Awerbuch et al. 94a] Baruch Awerbuch, Yair Bartal, Amos Fiat, and Adi Rosén. “Competitive Non-Preemptive Call Control.” In *Proc. 5th Symp. on Discrete Algorithms*, pp. 312–320, Philadelphia, PA: SIAM, 1994.
- [Awerbuch et al. 94b] Baruch Awerbuch, Rainer Gawlick, Tom Leighton, and Yuval Rabani. “On-Line Admission Control and Circuit Routing for High Performance Computing and Communication.” In *Proc. 35th Symp. Foundations of Computer Science*, pp. 412–423, Los Alamitos, CA: IEEE Press, 1994.
- [Blum et al. 01] Avrim Blum, Adam Kalai, and Jon Kleinberg. “Admission Control to Minimize Rejections.” In *Proceedings of WADS 2001*, Lecture Notes in Computer Science 2125, pp. 155–164, Berlin: Springer-Verlag, 2001.
- [Garay and Gopal 92] Juan A. Garay and I. S. Gopal. “Call Preemption in Communications Networks.” In *Proc. INFOCOM '92*, pp. 1043–1050, Los Alamitos, CA: IEEE Press, 1992.
- [Garay et al. 93] Juan A. Garay, I. S. Gopal, Shay Kutten, Yishay Mansour, and M. Yung. “Efficient On-Line Call Control Algorithms.” In *Proc. 2nd Israel Symp. on Theory of Computing and Systems*, pp. 285–293, Los Alamitos, CA: IEEE Press, 1993.

- [Kamath et al. 96] A. Kamath, O. Palmon, and Serge Plotkin. “Routing and Admission Control in General Topology Networks with Poisson arrivals.” In *Proc. 7th Symp. on Discrete Algorithms*, pp. 269–278, Philadelphia, PA: SIAM, 1996.
- [Leonardi 98] S. Leonardi. “On-Line Network Routing.” In *On-Line Algorithms*, edited by A. Fiat and G. Woeginger. Berlin: Springer-Verlag, 1998.
- [Leonardi et al. 98] Stefano Leonardi, Alberto Marchetti-Spaccamela, A. Presciutti, and Adi Rosèn. “On-Line Randomized Call-Control Revisited.” In *Proc. 9th Symp. on Discrete Algorithms*, pp. 323–332, Philadelphia, PA: SIAM, 1998.
- [Plotkin 95] Serge Plotkin. “Competitive Routing in ATM Networks.” *IEEE J. Selected Areas in Communications* (1995), 1128–1136.

---

Avrim Blum, Carnegie Mellon University, Pittsburgh PA 15213 (avrim@cs.cmu.edu)

Adam Kalai, Massachusetts Institute of Technology, Cambridge MA 02139  
(akalai@mit.edu)

Jon Kleinberg, Cornell University, Ithaca NY 14853 (kleinber@cs.cornell.edu)

Received December 13, 2002; accepted March 26, 2003.