Routing and Admission Control in General Topology Networks with Poisson Arrivals

Anil Kamath[*] *Omri Palmon*[†] *Serge Plotkin*[‡]

Abstract

Emerging high speed networks will carry traffic for services such as video-on-demand and video teleconferencing – that require resource reservation along the path on which the traffic is sent. High bandwidth-delay product of these networks prevents circuit rerouting, i.e. once a circuit is routed on a certain path, the bandwidth taken by this circuit remains unavailable for the duration (holding time) of this circuit. As a result, such networks will need effective *routing* and *admission control* strategies.

Recently developed online routing and admission control strategies have logarithmic competitive ratios with respect to the *admission ratio* (the fraction of admitted circuits). Such guarantees on performance are rather weak in the most interesting case where the rejection ratio of the optimum algorithm is very small or even 0. Unfortunately, these guarantees can not be improved in the context of the considered models, making it impossible to use these models to identify algorithms that are going to perform well in practice.

In this paper we develop routing and admission control strategies for a more realistic model, where the requests for virtual circuits between any two points arrive according to a Poisson process and where the circuit holding times are exponentially distributed. Our model is close to the one that was developed to analyse and tune the (currently used) strategies for managing traffic in long-distance telephone networks. We strengthen this model by assuming that the rates of the Poisson processes (the "traffic matrix") are unknown to the algorithm and are chosen by the adversary.

Our strategy is competitive with respect to the expected *rejection ratio*. More precisely, it achieves expected rejection ratio of at most $R^* + \epsilon$, where R^* is the optimum expected rejection ratio. The expectations are taken over the distribution of the request sequences, and $\epsilon = O(\sqrt{r \log n})$, where r is the maximum fraction of an edge bandwidth that can be requested by a single circuit.

Our result should be viewed in context of the previous competitive routing and admission control strategies that require $r \leq 1/\log n$, but are not able to formally analyse the (intuitively clear) relation between r and the performance achievable in realistic situations.

1 Introduction

1.1 Overview

In order to provide quality of service guarantees, the new high-speed networks will allocate resources in terms of virtual circuits. In particular, creating a virtual circuit will require reservation of

[*]Department of Computer Science, Stanford University, Stanford, CA 94305. E-Mail: kamath@cs.stanford.edu.
Supported by NSF Grant CCR-9304971.
[†]Department of Computer Science, Stanford University, Stanford, CA 94305. E-Mail: omri@cs.stanford.edu.
Supported by NSF Grant CCR-9304971.
[‡]Department of Computer Science, Stanford University, Stanford, CA 94305. E-Mail: plotkin@cs.stanford.edu.
Supported by ARO Grant DAAH04-95-1-0121, NSF Grant CCR-9304971 and Terman Fellowship.

bandwidth on some path between the endpoints of the connection. Admission control algorithms are needed since network resources are limited.

Due to high bandwidth-delay product of the future high-speed networks, rerouting of circuits is not supported. In other words, once the circuit is routed along some path, the resources along this path are unavailable for the duration (or "holding time") of this circuit. The task of the *routing* and *admission control* strategy is to decide which circuits should be rejected vs. those that should be accepted, and to choose the paths for the accepted circuits.

The problem of admission control and routing in general topology networks was considered in the context of *competitive analysis* [35, 18, 11, 25] in [5, 6, 7]. Roughly speaking, the statement that a particular strategy has competitive ratio α means that its performance is at most a factor of $1/\alpha$ of the performance of the best possible offline algorithm. The congestion-minimization model was considered in [5, 10], where an algorithm was proposed that can route all of the requests using $O(\log nT)$-factor more capacity than the adversary, where n is the number of nodes in the network and T is the ratio of the maximum to the minimum holding time of the circuit. The throughput model was considered in [6], where it was shown how to modify the algorithm in [5] to make it accept at least $O(\log LT)$ fraction of the circuits accepted by the adversary (using the same capacity), where L is the maximum allowable number of edges (hops) used by a circuit.[1]

The main disadvantage of the above competitive strategies is that they provide very weak guarantees on performance. Although one can argue that the value of L is relatively small, the value of T can easily be in the hundreds or thousands. For example, consider a very small network where L=4 and T=128. If a request sequence can be completely satisfied by the adversary, the algorithm in [6] guarantees to accept at least 3% of the circuits, which is not very useful from the practical point of view. Unfortunately, the lower bounds in [5, 6] imply that one can not do better in this model. In fact, if the holding time of a circuit is unknown at the time the request to route this circuit is issued, it is possible to prove even stronger lower bounds [9, 24]. This lead [14] to consider an alternate model, where the arrival times of requests are arbitrary, but the holding times are distributed exponentially. For this model, they showed a strategy that routes expected $O(1/\log L)$ fraction of circuits routed by the optimum algorithm that has a complete knowledge of request arrivals. This bound is still not useful from the practical point of view.

Another problem with the competitive approaches mentioned above is that they can not be used to analyze the intuitively obvious fact that the performance increases if we decrease r, the maximum fraction of link capacity that can be requested by a single circuit. All of the above strategies require that $r \leq 1/\log n$, but do not suggest how to improve the performance when r is much smaller.[2]

There is a large body of work on practical routing and admission control strategies for long-distance telephone networks (POTS). Examples include reservation based strategies such as RTNR used by AT&T [3] and DNHR [2, 4, 1] used by British Telecom (see also [19, 28, 29, 27, 26, 16]). The analyses of these strategies assume that the requests arrive as a Poisson process, and the holding times are distributed exponentially. Moreover, they assume that the underlying graph is complete and that the traffic matrix is fixed and uniform, i.e. same rate of arrival for circuits between any pair of points. It is important to mention that all of the above papers assume that the processes describing utilization of different edges are independent. Such assumption might be justified when the underlying topology is a complete graph where almost all of the circuits are routed on a single edge (which is the case in POTS), but is clearly wrong for general topology networks.

[1]The case of a single-edge network was considered in [13], and line network was considered in [12]; see [31] for the survey of competitive routing strategies.

[2]Algorithms for special topologies that do not rely on this assumption are given in [8, 21].

1.2 Our results

In this paper we consider the case where the requests arrive according to a Poisson distribution and where the holding times are exponentially distributed. We bound the expected performance of our algorithm in terms of the expected performance of the offline algorithm, where the expectation is taken over the distribution on the inputs, and where the offline algorithm has a complete a priori knowledge of all the requests and all the termination times.

Let $\epsilon = \sqrt{r \log n}$, where r the maximum fraction of an edge bandwidth that can be requested by a single circuit. If an offline algorithm can achieve expected rejection ratio R^*, then our algorithm achieves expected rejection ratio of at most $R^* + O(\epsilon)$.

In contrast to the previously proposed strategies in the Poisson-arrivals model [19, 22, 30], we do not assume that the processes describing utilization of different edges are independent. Moreover, contrary to these papers, we assume that the traffic matrix (rates of arrivals between source/destination pairs) is *unknown* and chosen by the adversary.

Our algorithm is similar to the algorithms in [5, 10, 6] in that it routes along a shortest path, where the length is an exponential function of edge congestion. Rejections occur if the length of the shortest path exceeds a threshold. As opposed to the previous approaches that used a static threshold, our algorithm uses a *dynamic threshold* that depends on the previous rejections.

The model and the analysis technique proposed in this paper show how to take advantage of the situations where r, the maximum fraction of an edge bandwidth that can be requested by a single circuit, is very small, i.e. $r \ll 1/\log n$. (The fact that small r leads to better performance was empirically observed in a simulation study in [15].) This is in contrast to the previous approaches which required $r \leq 1/\log n$, but could not improve performance for the cases where r is smaller.

The routing and admission control strategies developed in this paper, draw heavily on the approach developed in the context of approximation algorithms for the multicommodity flow problem [33, 20, 23] and in particular, minimum-cost multicommodity flow problem [32, 17]. Roughly speaking, these algorithms are based on the following idea: define a potential function such that if it is sufficiently small, then the current solution is close to optimum. Repeatedly compute the gradient of the potential function and use it to *reroute* some of the flow. The heart of these algorithms is the proof that the proper choice of the potential function together with the rerouting step cause reduction in the potential function. In our case, we can not use rerouting, and hence it is impossible to apply these algorithms directly. Instead, we show that our routing strategy *maintains* the expectation of the potential function below a certain threshold. Because of the choice of the potential function, this implies high-probability bounds on the congestion.

It is interesting to note that adding a simple scaling to our strategy (such as in [23]), results in an algorithm that can be used to compute approximate multicommodity flow. The resulting algorithm is somewhat similar to the one in [36], but slower than the algorithms in [23, 32, 17].

It seems that a possible alternative to our algorithm is to learn the traffic matrix and then route according to a fractional multicommodity flow solution. In practice it is possible that the traffic matrix will change before sufficiently good statistical estimations can be computed. In contrast to methods that relay on availability of statistical information [22, 30, 34], our algorithm achieves provably good performance even if the traffic matrix changes in discrete intervals. In particular, in the congestion minimization case, our algorithm achieves expected congestion of $1 + O(\epsilon)$ larger than the expected congestion achieved by the offline algorithm on the worst one of the matrices that were used to generate the traffic.

Using a transformation similar to the one in [5], one can modify our algorithm to solve the online load-balancing problem for unrelated machine scheduling. If job durations are distributed exponentially and jobs arrivals are Poisson with unknown rates (different rates for different job types), we can achieve expected load that is within $1 + O(\epsilon)$ of the expected load of the offline algorithm, where the expectation is over the input sequences, with the same ϵ as above.

2 Model and Definitions

The network is represented as a capacitated graph $G(V, E, u)$ (the results apply both to directed and undirected cases), where $u(e)$ represents the capacity of the edge $e \in E$. Request for a virtual circuit is a tuple (v, w, r, ρ) where nodes v and w are the source and destination of the request, r is the requested bandwidth, and ρ is the "profit" per time unit that we get if the request is routed (accepted).

When presented with request (v, w, r, ρ), the goal of the admission control and routing strategy is to decide whether to accept or to reject the request and, if the decision is to accept, to reserve bandwidth r along some path between the endpoints v and w. This bandwidth remains reserved for the duration (holding time) of this circuit. Congestion on edge e at time t, i.e. the fraction of the bandwidth of edge e reserved at time t, is denoted by $\Lambda_t(e)$.

We will consider the model where we have k types of circuits. Type i is associated with source/destination pair v_i, w_i and bandwidth r_i. We assume that the arrival of requests for type i circuits is a Poisson process with *unknown* rate λ_i. Accepting type i circuit will generate profit ρ_i per time unit. We will assume that the holding times are exponential with mean T.

In this paper we concentrate on two related models. In the *congestion-minimization* model, the routing strategy is required to accept all of the requests. The goal is to minimize the maximum (over all edges) congestion Λ. In the *throughput-maximization* model, the routing strategy is allowed to reject some of the requests. The goal in this case is to maximize the total profit associated with the accepted requests, while not exceeding the available capacity.

In both models, we compare performance of our algorithms to the performance of an off-line algorithm that has a priori knowledge of all the requests and termination times. We prove that, with very high probability, the performance of our algorithm (measured either by maximal congestion or by profit), is within a small factor from the average performance of the off-line algorithm. The probability and average are taken with respect to the distribution of the inputs.

Let $f_t^i(e)$ denote the total bandwidth allocated on edge e at time t for type i circuits. We will use f_t^i to denote the vector $(f_t^i(e_1), f_t^i(e_2), \dots)$ and $f_t = \sum_i f_t^i$. It will help to view f_t as a multiflow where there is a commodity for each request type, and the demand for commodity i is determined by the total bandwidth of circuits of type i that are alive at time t.

Congestion on edge e at time t is defined by

$$\Lambda_t(e) = \Lambda(f_t(e)) = \sum_i f_t^i(e)/u(e)$$

where $u(e)$ is the bandwidth of this edge. The overall congestion of f_t is denoted by $\Lambda(f_t) = \max_e\{\Lambda(f_t(e))\}$.

In order to relate the performance of our algorithm to the performance of an off-line algorithm we consider a multicommodity flow f^* associated with the traffic matrix. In the context of the congestion model, f^* denotes a solution to the concurrent flow problem [33], with demands $d_i = \lambda_i r_i T$ (the average amount of bandwidth required for circuits of type i). In other words, f^* is a flow that satisfies demands $\lambda_i r_i T$, and minimizes congestion $\Lambda^* = \Lambda(f^*) = \max_e\{f^*(e)/u(e)\}$.

In the context of the throughput-maximization model, we will use f^* to denote the solution to the max-sum multicommodity flow, i.e. f^* will be a flow that satisfies demands $d'_i \le d_i$, satisfies capacity constraints, and maximizes the profit due to the satisfied demands $\sum d'_i \rho_i$.

3 Congestion-Minimization Strategy

In this section we describe a congestion-minimization routing strategy. We show that this strategy ensures that during overwhelming proportion of the time, the maximum congestion stays within $(1+\epsilon)$ of the expected congestion achieved by an optimum off-line algorithm. This will be done by relating expected congestion to Λ^*, the lower bound on the expected congestion given by a solution of the corresponding concurrent flow problem. We will assume that the capacities are scaled such that $\Lambda^* \leq 1$.

3.1 Online Algorithm

The algorithm maintains the current bandwidth assigned to each edge, i.e. the current congestion $\Lambda_t(e)$. A request from source node v to destination w that arrives at time t is routed along a shortest path with respect to edge-costs given by:

$$\text{cost}(e) = \frac{1}{u(e)} a^{\Lambda_t(e)}$$

The parameter a is defined by

$$a = (p^{-1}m)^{1/\epsilon} \tag{1}$$

where m is the number of edges. The values of p and ϵ depend on the optimization criteria and will be discussed later. Roughly speaking, we will show that p represents the fraction of time when the congestion exceeds Λ^* by more than a factor of $(1+O(\epsilon))$. We assume that the requested bandwidth r_i satisfies the following *granularity* condition for each edge e it can be routed on:

$$\frac{r_i}{u(e)} \leq \frac{\epsilon^2}{\log p^{-1}m} \tag{2}$$

Note that, for constant ϵ, this condition is essentially the same as the granularity condition required in [6].

3.2 Bounding Offline Performance

Theorem 3.1 *Let $\mathcal{M}$ be an optimum offline algorithm that has a priori knowledge of all the requests and the holding times. Define a random variable $\Lambda_t^{\mathcal{M}}$ to be the maximal congestion achieved by $\mathcal{M}$ at time t; let f^* and $\Lambda(f^*)$ be defined as in Section 2. Then $\mathbb{E}\{\Lambda_t^{\mathcal{M}}\} \geq \Lambda^*$, where the expectation is taken with respect to the distribution of arrivals and terminations of circuits.*

Proof: Define a vector of random variables $(f_t^{\mathcal{M}}(e;i))_{i,e}$ as the bandwidth of edge e assigned by $\mathcal{M}$ to circuits of type i at time t. Define the random variables

$$\begin{aligned} \Lambda_t^{\mathcal{M}}(e) &= \frac{1}{u(e)} \sum_i f_t^{\mathcal{M}}(e;i) \\ \Lambda_t^{\mathcal{M}} &= \max_e \Lambda_t^{\mathcal{M}}(e) \end{aligned}$$

Observe that $\mathbb{E}\{f_t^{\mathcal{M}}(e;i)\}$ is a multicommodity flow that satisfies demands $d_i = \lambda_i r_i T$ between v_i and w_i. Its congestion at e is given by $\mathbb{E}\{\Lambda_t^{\mathcal{M}}(e)\}$. Since Λ^* is the optimum congestion, we have $\Lambda^* \leq \max_e \mathbb{E}\{\Lambda_t^{\mathcal{M}}(e)\}$. It remains to observe that $\max_e \mathbb{E}\{\Lambda_t^{\mathcal{M}}(e)\} \leq \mathbb{E}\{\max_e \Lambda_t^{\mathcal{M}}(e)\}$, which is true since for every e we have $\mathbb{E}\{\Lambda_t^{\mathcal{M}}(e)\} \leq \mathbb{E}\{\max_e \Lambda_t^{\mathcal{M}}(e)\}$. ∎

3.3 Analysis

Given any multiflow f, we define the following potential function:

$$\Psi(f) = \sum_e a^{\Lambda(f(e))}$$

where $a = (p^{-1}m)^{1/\epsilon}$. The performance bound on the routing algorithm will be proved by a sequence of lemmas. First we observe that if $\Psi(f_t)$ is not much larger than $\Psi(f^*)$, then $\Lambda(f_t)$ is very close to $\Lambda(f^*)$, i.e. very close to optimum. The heart of the proof is the claim that as long as $\Psi(f_t)$ is larger than $\Psi((1+3\epsilon)f^*)$, the derivative with respect to time of the expected value of $\Psi(f_t)$ is negative. From this we conclude that the expectation of $\Psi(f_t)$ stays below $\Psi((1+3\epsilon)f^*)$, which implies that, with high probability, $\Lambda(f_t)$ does not exceed $(1+O(\epsilon))\Lambda(f^*)$.
The following lemma shows that small value of Ψ implies small congestion.

Lemma 3.2 *Let f_1, f_2 be two multiflows. Assume that $\Psi(f_1) \leq \alpha\Psi(f_2)$. Then $\Lambda(f_1) \leq \frac{\log \alpha m}{\log a} + \Lambda(f_2)$.*

Proof: From the definition and Ψ we know that $\Psi(f_1) \geq a^{\Lambda(f_1)}$ and $\Psi(f_2) \leq ma^{\Lambda(f_2)}$. Hence $a^{\Lambda(f_1)} \leq \alpha m a^{\Lambda(f_2)}$, and the lemma follows. ▮

We will fix some time t_0 and analyze the behavior of the potential function at t_0. We will also work with $t > t_0$ and denote $\Delta t = t - t_0$. Denote by $\Delta_f^{t,t_0}(e)$ the random process associated with the change in the congestions on edge e in the interval $[t_0, t]$: $\Delta_f^{t,t_0}(e) = \Lambda_t(e) - \Lambda_{t_0}(e)$. Observe that $\Delta_f^{t,t_0}(e)$ is determined by two types of events: arrival of new circuits, and departures of existing ones. Instead of using $\Delta_f^{t,t_0}(e)$, we will consider a simpler process $\bar{\Delta}_f^{t,t_0}(e)$ that will provide a sufficiently good approximation to $\Delta_f^{t,t_0}(e)$ around t_0. Towards this end, we will modify the events associated with $\Delta_f^{t,t_0}(e)$ in two different ways.

First, consider the circuit terminations that affect $\Delta_f^{t,t_0}(e)$. Termination of a circuit of type i that was using edge e, causes $\Delta_f^{t,t_0}(e)$ to be decremented by r_i. To compute $\bar{\Delta}_f^{t,t_0}(e)$, we associate each circuit that exists at time t_0 with a Poisson process of rate $1/T$. Each firing of such process associated with a circuit of type i causes $\bar{\Delta}_f^{t,t_0}(e)$ to be decremented by r_i. Observe that because the holding times are distributed exponentially with mean T, the first event (and only the first event) of this process corresponds to the real termination event of the corresponding circuit.

The second difference between $\Delta_f^{t,t_0}(e)$ and $\bar{\Delta}_f^{t,t_0}(e)$ is due to the different way we compute the increments. The increments in $\Delta_f^{t,t_0}(e)$ are due to routing a circuit on edge e. A circuit is routed on the shortest paths with respect to the costs computed at its arrival time. Instead, the increments in $\bar{\Delta}_f^{t,t_0}(e)$ will be computed as if all the circuits arriving in $[t_0, t]$ are routed on shortest paths with respect to the costs at t_0.

The exact distribution for $\bar{\Delta}_f^{t,t_0}(e)$ can now be written as

$$\frac{1}{u(e)} \sum_{i \leq k} r_i P\left(\frac{\bar{f}^i(e)}{r_i}\Delta t\right) - r_i P\left(\frac{f_{t_0}^i(e)}{r_i}\frac{1}{T}\Delta t\right)$$

where $\bar{f}$ is an uncapacitated fractional multicommodity flow that satisfies the demands $\lambda_i r_i$ and minimizes the cost $\sum_e \bar{f}(e) a^{\Lambda_{t_0}(e)}/u(e)$, and $P(\lambda)$ is a Poisson process with rate λ.

Every event (e.g. routing of a new circuit or termination of an existing one) causes an increment or a decrement in $\Delta_f^{t,t_0}(e)$ and $\bar{\Delta}_f^{t,t_0}(e)$. Note that some of the events (those that are associated with the modified termination processes) affect $\bar{\Delta}_f^{t,t_0}(e)$ only. Observe that Δ_f^{t,t_0} and $\bar{\Delta}_f^{t,t_0}$ are different only if at least 2 events of any type occured in the $[t_0, t]$ interval. This observation is the basis of the following Lemma.

Lemma 3.3
$$\left|\mathbb{E}\left\{a^{\Delta_f^{t,t_0}(e)}\right\} - \mathbb{E}\left\{a^{\bar\Delta_f^{t,t_0}(e)}\right\}\right| \le Q(\Delta t)^2$$
where Q is independent of t, Δt, and the state at t_0.

Proof: Let X be a random variable which counts the number of events that change $\bar\Delta_f^{t,t_0}(e)$. Note that X does not count all of the events. Conditioning on X, we get

$$\left|\mathbb{E}\left\{a^{\Delta_f^{t,t_0}(e)}\right\} - \mathbb{E}\left\{a^{\bar\Delta_f^{t,t_0}(e)}\right\}\right| \le \mathbb{E}\left\{\left|a^{\Delta_f^{t,t_0}(e)} - a^{\bar\Delta_f^{t,t_0}(e)}\right|\right\}$$

$$\le \sum_{j=0}^{\infty} \mathrm{P}rob(X=j) \max_{X=j} \left|a^{\Delta_f^{t,t_0}(e)} - a^{\bar\Delta_f^{t,t_0}(e)}\right|$$

By definition of X, we have
$$\max_{X=j} \left|a^{\Delta_f^{t,t_0}(e)} - a^{\bar\Delta_f^{t,t_0}(e)}\right| \le a^{\frac{rj}{u(e)}}$$
where $r = \max_i r_i$. If $X = 0$ then $\Delta_f^{t,t_0}(e) = \bar\Delta_f^{t,t_0}(e)$. If $X = 1$, we have $\Delta_f^{t,t_0}(e) = \bar\Delta_f^{t,t_0}(e)$ as well, unless the event that is counted by X is an arrival of a new circuit that departs before time t. Let A denote this combination, i.e. arrival and termination of a single circuit in the (t_0, t) interval. Thus

$$\left|\mathbb{E}\left\{a^{\Delta_f^{t,t_0}(e)}\right\} - \mathbb{E}\left\{a^{\bar\Delta_f^{t,t_0}(e)}\right\}\right| \le \sum_{j=2}^{\infty} \mathrm{P}rob(X=j) a^{\frac{rj}{u(e)}} + \mathrm{P}rob(A) a^{\frac{r}{u(e)}}$$

Recall that $\Delta t = t - t_0$. Let l denote the number of circuits on e at time t_0. Observe that X is a Poisson process which counts two types of events: arrivals of new circuits, which is a Poisson process with average $\sum \lambda_i \Delta t$, and events that decrement $\bar\Delta_f^{t,t_0}(e)$, which is a Poisson process with average $l\Delta t/T$. Denote the rate of increase in X by $\bar\lambda = \sum \lambda_i + l/T$; the average of X is $\bar\lambda \Delta t$. Then

$$\sum_{j=2}^{\infty} \mathrm{P}rob(X=j) a^{\frac{rj}{u(e)}} = \sum_{j=2}^{\infty} e^{-\bar\lambda\Delta t} \frac{(\bar\lambda \Delta t)^j}{j!} a^{\frac{rj}{u(e)}} =$$

$$e^{-\bar\lambda\Delta t}\left(e^{\bar\lambda\Delta t a^{\frac{r}{u(e)}}} - 1 - \bar\lambda\Delta t a^{\frac{r}{u(e)}}\right) \le e^{-\bar\lambda\Delta t}(\bar\lambda\Delta t a^{\frac{r}{u(e)}})^2 e^{\bar\lambda\Delta t a^{\frac{r}{u(e)}}}$$

where the last inequality follows from the fact that for all $x > 0$, we have $e^x \le 1 + x + x^2 e^x$. The probability of the event A can be estimated by

$$\mathrm{P}rob(A) \le \sum_i \lambda_i \Delta t (1 - e^{-\Delta t/T})$$

Combining $\Delta t < 1$ with these estimates implies the claim of the lemma. ∎

From the above lemma, we see that $\bar\Delta_f^{t,t_0}(e)$ is a good approximation to $\Delta_f^{t,t_0}(e)$:

Corollary 3.4 $$\frac{\partial}{\partial t}\, \mathbb{E}\left\{a^{\Delta_f^{t,t_0}(e)}\right\}\Big|_{t=t_0} = \frac{\partial}{\partial t}\, \mathbb{E}\left\{a^{\bar\Delta_f^{t,t_0}(e)}\right\}\Big|_{t=t_0}$$

Roughly speaking, the following theorem states that if the expectation of $\Psi(f_{t_0})$ exceeds $\Psi((1 + 3\epsilon f^*))$, then its derivative with respect to time is negative.

Theorem 3.5 *Assume that granularity condition (2) is satisfied for some $\epsilon < 1/3$. Let f_t be the flow generated by the algorithm at t. Then*

$$\frac{\partial}{\partial t}\,\mathbb{E}\{\Psi(f_t)\}|_{t=t_0} \le \frac{1}{T}(1-\epsilon)\left(\Psi((1+3\epsilon)f^*) - \Psi(f_{t_0})\right),$$

where f^* is a flow that satisfies the demands $d_i = \lambda_i r_i T$.

Proof: By definition of $\Delta_f^{t,t_0}(e)$, we have

$$\mathbb{E}\{\Psi(f_t)\} = \sum_e \mathbb{E}\left\{a^{\Lambda_t(e)}\right\} = \sum_e a^{\Lambda_{t_0}(e)}\mathbb{E}\left\{a^{\Delta_f^{t,t_0}(e)}\right\}$$

Using Corollary 3.4, we get

$$\frac{\partial}{\partial t}\,\mathbb{E}\{\Psi(f_t)\}|_{t=t_0} = \sum_e a^{\Lambda_{t_0}(e)}\frac{\partial}{\partial t}\,\mathbb{E}\left\{a^{\bar{\Delta}_f^{t,t_0}(e)}\right\}\Big|_{t=t_0} \tag{3}$$

To calculate $\mathbb{E}\left\{a^{\bar{\Delta}_f^{t,t_0}(e)}\right\}$ we will use the fact that the moment generating function of random variable X that is distributed as Poisson with rate λ, is equal to $M\left(X;\theta\right) = \mathbb{E}\{e^{\theta X}\} = \exp\left(\lambda(e^\theta - 1)\right)$. We will denote by $\bar{P}_i$ the Poisson variable that counts the number of arrival of new circuits of type i, and P_i will denote the Poisson variable that counts the number of terminations of circuits of type i. We know that $\bar{P}_i$ has rate $\frac{\bar{f}^i(e)}{r_i}\Delta t$ and P_i has rate $\frac{f^i_{t_0}(e)}{r_i}\frac{1}{T}\Delta t$. Taking into account that the arrivals and the terminations are independent, we get:

$$\begin{aligned}
\mathbb{E}\left\{a^{\bar{\Delta}_f^{t,t_0}(e)}\right\} &= M\left(\bar{\Delta}_f^{t,t_0}(e);\log a\right) \\
&= \prod_{i\le k} M\left(r_i\bar{P}_i - r_iP_i;\frac{1}{u(e)}\log a\right) \\
&= \prod_{i\le k} M\left(\bar{P}_i;\frac{r_i}{u(e)}\log a\right)\cdot M\left(P_i;-\frac{r_i}{u(e)}\log a\right) \\
&= \exp\left(\sum_{i\le k}\frac{\bar{f}^i(e)}{r_i}\Delta t(e^{\log a\frac{r_i}{u(e)}}-1)+\frac{f^i_{t_0}(e)}{r_i}\frac{1}{T}\Delta t(e^{-\log a\frac{r_i}{u(e)}}-1)\right)
\end{aligned}$$

Hence

$$\frac{\partial}{\partial t}\,\mathbb{E}\left\{a^{\bar{\Delta}_f^{t,t_0}(e)}\right\}\Big|_{t=t_0} = \sum_{i\le k}\frac{\bar{f}^i(e)}{r_i}(e^{\frac{r_i\log a}{u(e)}}-1)+\frac{f^i_{t_0}(e)}{r_i}\frac{1}{T}(e^{\frac{-r_i\log a}{u(e)}}-1)$$

Since $\log a\frac{r_i}{u(e)} \le \epsilon \le 1$ and $e^x \le 1 + x + x^2$ for $|x| \le 1$, we have

$$\frac{\partial}{\partial t}\,\mathbb{E}\left\{a^{\bar{\Delta}_f^{t,t_0}(e)}\right\}\Big|_{t=t_0} \le$$

$$\sum_{i\le k}\frac{\bar{f}^i(e)}{r_i}\frac{r_i\log a}{u(e)} - \frac{f^i_{t_0}(e)}{r_i}\frac{1}{T}\frac{r_i\log a}{u(e)} + \frac{\bar{f}^i(e)}{r_i}\left(\frac{r_i\log a}{u(e)}\right)^2 + \frac{f^i_{t_0}(e)}{r_i}\frac{1}{T}\left(\frac{r_i\log a}{u(e)}\right)^2 \le$$

$$\log a\frac{1}{T}\left[(1+\epsilon)\Lambda(T\bar{f}(e)) - (1-\epsilon)\Lambda_{t_0}(e)\right]$$

Combining this estimate with (3), we get

$$\frac{\partial}{\partial t}\,\mathbb{E}\{\Psi(f_t)\}|_{t=t_0} \le \frac{\log a}{T}\sum_e a^{\Lambda_{t_0}(e)}\,[(1+\epsilon)\Lambda(T\bar{f}(e)) - (1-\epsilon)\Lambda_{t_0}(e)]$$

Recall that $\bar{f}$ satisfies demands $\lambda_i r_i$ and flows along the *shortest paths* with respect to the costs $a^{\Lambda_{t_0}(e)}/u(e)$. The fact that f^* satisfies the demands $\lambda_i r_i T$ implies that

$$\sum_e \Lambda(T\bar{f}(e))a^{\Lambda_{t_0}(e)} \le \sum_e \Lambda(f^*(e))a^{\Lambda_{t_0}(e)}$$

Hence,

$$\begin{aligned}
\frac{\partial}{\partial t}\,\mathbb{E}\{\Psi(f_t)\}|_{t=t_0} &\le \log a\frac{1}{T}\sum_e a^{\Lambda_{t_0}(e)}\,[(1+\epsilon)\Lambda(f^*(e)) - (1-\epsilon)\Lambda_{t_0}(e)] \\
&= \frac{1}{T}\frac{\partial}{\partial t}\,\Psi(f_{t_0} + t(f^* - f_{t_0} + \epsilon(f^* + f_{t_0})))|_{t=t_0} \\
&= \frac{1}{T}\frac{\partial}{\partial t}\,\Psi(f_{t_0} + (1-\epsilon)t(\frac{1+\epsilon}{1-\epsilon}f^* - f_{t_0}))\Big|_{t=t_0} \\
&= \frac{1}{T}(1-\epsilon)\frac{\partial}{\partial t}\,\Psi\left(f_{t_0} + t(\frac{1+\epsilon}{1-\epsilon}f^* - f_{t_0})\right)\Big|_{t=t_0}
\end{aligned}$$

Since Ψ is convex, the last expression is bounded by

$$\frac{1}{T}(1-\epsilon)\left(\Psi\left(\frac{1+\epsilon}{1-\epsilon}f^*\right) - \Psi(f_{t_0})\right) \le \frac{1}{T}(1-\epsilon)\left[\Psi\left((1+3\epsilon)f^*\right) - \Psi(f_{t_0})\right]$$

which proves the theorem. ∎

The previous theorem provides bounds on the derivative with respect to time of expectation of $\Psi(f_t)$, where the expectation is taken over all possible arrival/departure events after time t_0. In particular, it implies that if $\Psi(f_{t_0}) < \Psi((1+3\epsilon)f^*)$, then this derivative is negative. We would like to infer from this that the expectation of $\Psi(f_t)$ never exceeds $\Psi((1+3\epsilon)f^*)$. Towards this end, we need to strengthen the above theorem and show a similar bound on the derivative of the expectation, where the expectation is taken over *all of the events up to* t, and not only the events that occured after time t_0.

Theorem 3.6 *Assume that at time $t = 0$ the system was in a state where $\Lambda_e(f_0) \le \Lambda_e(f^*)$, and at time $t_0 \ge 0$ it reached a state defined by probability measure μ. Let ν denote the probability measure of the events that occur after time t_0. Denote by f_t the (random) flow that is generated by the algorithm after time t elapsed, where f_t is determined both by μ and ν. Then*

$$\frac{\partial}{\partial t}\,\mathbb{E}_{\mu,\nu}\left\{\Psi(f_t)\right\}|_{t=t_0} \le$$
$$\frac{1}{T}(1-\epsilon)\left(\Psi\left((1+3\epsilon)f^*\right) - \mathbb{E}_\mu\left\{\Psi(f_{t_0})\right\}\right)$$

Where f^* any flow that satisfies the demands $d_i = \lambda_i r_i T$.

Proof:

$$\begin{aligned}
\frac{\partial}{\partial t}\,\mathbb{E}_{\mu,\nu}\left\{\Psi(f_t)\right\}|_{t=t_0} &= \lim_{t\to t_0}\frac{\mathbb{E}_{\mu,\nu}\left\{\Psi(f_t)\right\} - \mathbb{E}_{\mu,\nu}\left\{\Psi(f_{t_0})\right\}}{\Delta t} \\
&= \lim_{t\to t_0}\frac{\mathbb{E}_\mu\left\{\mathbb{E}_\nu\left\{\Psi(f_t)\right\}\right\} - \mathbb{E}_\mu\left\{\Psi(f_{t_0})\right\}}{\Delta t} \\
&= \lim_{t\to t_0}\mathbb{E}_\mu\left\{\frac{\mathbb{E}_\nu\left\{\Psi(f_t)\right\} - \Psi(f_{t_0})}{\Delta t}\right\}
\end{aligned}$$

At this point we would like to exchange the order of the limit and the expectation. To justify replacing the limit and expectation we first note that it is enough to prove that for every edge e

$$\lim_{t\to t_0}\mathbb{E}_\mu\left\{a^{\Lambda_{t_0}(e)}\frac{\mathbb{E}_\nu\left\{a^{\Delta_f^{t,t_0}(e)}\right\}-1}{\Delta t}\right\}=\mathbb{E}_\mu\left\{\lim_{t\to t_0}a^{\Lambda_{t_0}(e)}\frac{\mathbb{E}_\nu\left\{a^{\Delta_f^{t,t_0}(e)}\right\}-1}{\Delta t}\right\}$$

Instead of proving the equality directly we can prove the same equality for $\bar{\Delta}_f^{t,t_0}(e)$ instead of $\Delta_f^{t,t_0}(e)$ and then use the bound in Lemma 3.3. By Lebesgue's bounded convergence theorem it is enough to show that there is a random variable g, defined on the same probability space as μ, such that for all t in some neighborhood of t_0

$$a^{\Lambda_{t_0}(e)}\frac{\mathbb{E}_\nu\left\{a^{\bar{\Delta}_f^{t,t_0}(e)}-1\right\}}{\Delta t}<g\quad\text{point wise}$$

and $\mathbb{E}_\mu\{g\}<\infty$. The existence of such g is proved in Lemma 3.7. The fact that g has finite integral is proved in 3.8.

Using the fact that the limit and the expectation can be exchanged we get that the original expression is equal to:

$$\begin{aligned}\lim_{t\to t_0}\mathbb{E}_\mu\left\{\frac{\mathbb{E}_\nu\{\Psi(f_t)\}-\Psi(f_{t_0})}{\Delta t}\right\} &= \mathbb{E}_\mu\left\{\lim_{t\to t_0}\frac{\mathbb{E}_\nu\{\Psi(f_t)\}-\Psi(f_{t_0})}{\Delta t}\right\}\\ &= \mathbb{E}_\mu\left\{\frac{\partial}{\partial t}\mathbb{E}_\nu\{\Psi(f_t)\}|_{t=t_0}\right\}\\ &\le \mathbb{E}_\mu\left\{\frac{1}{T}(1-\epsilon)\left(\Psi((1+3\epsilon)f^*)-\Psi(f_{t_0})\right)\right\}\\ &= \frac{1}{T}(1-\epsilon)\left(\Psi((1+3\epsilon)f^*)-\mathbb{E}_\mu\{\Psi(f_{t_0})\}\right)\end{aligned}$$

where the inequality is implied by Theorem 3.5. ∎

Lemma 3.7 *There exist constants b_1, b_2, b_3 such that for all $t\le 1$*

$$a^{\Lambda_{t_0}(e)}\frac{\mathbb{E}_\nu\left\{a^{\bar{\Delta}_f^{t,t_0}(e)}-1\right\}}{\Delta t}\le b_1(b_2^l+lb_3^l)$$

where l is the number of circuits in the system and where the constants depend on the initial state.

Proof: Let r be the maximal r_i. Let X be defined as in Lemma 3.3, i.e. the number events changing $\bar{\Delta}_f^{t,t_0}(e)$ in the interval $[t_0,t]$. X is a Poisson random variable with mean $\mathbb{E}\{X\}=\Delta t\left(\sum_i\lambda_i+l/T\right)=\Delta t\bar{X}$. *Note that X is defined differently for each initial state.* Define a random variable g_1 on product of the probability spaces of μ and ν as

$$g_1=\begin{cases}a^{X\frac{r}{u(e)}} & X>0\\ 0 & X=0\end{cases}$$

X was defined such that $\left|\bar{\Delta}_f^{t,t_0}(e)\right|\le X\frac{r}{u(e)}$ and hence $a^{\bar{\Delta}_f^{t,t_0}(e)}-1\le g_1$ point wise (since when

$X = 0$ we know that $a^{\bar{\Delta}_f^{t,t_0}(e)} = 1$). Estimating the expected value of g_1 gives

$$\begin{aligned}
\mathbb{E}_\nu\{g_1\} &= \sum_{j=1}^{\infty} a^{j\frac{r}{u(e)}} \mathrm{P}rob(X = j) \\
&= e^{-\Delta t \bar{X}} \sum_{j=1}^{\infty} a^{j\frac{r}{u(e)}} \frac{(\Delta t \bar{X})^j}{j!} \\
&= e^{-\Delta t \bar{X}} \left(e^{\Delta t \bar{X} a^{\frac{r}{u(e)}}} - 1 \right) \\
&\leq e^{-\Delta t \bar{X}} \Delta t \bar{X} a^{\frac{r}{u(e)}} e^{\Delta t \bar{X} a^{\frac{r}{u(e)}}}
\end{aligned}$$

Where the last inequality follows from the fact that $e^x \leq 1 + xe^x$ for $x \geq 0$. Using the last estimate and the definition of l gives

$$a^{\Lambda_{t_0}(e)} \frac{\mathbb{E}_\nu \left\{ a^{\bar{\Delta}_f^{t,t_0}(e)} - 1 \right\}}{\Delta t} \leq a^{\frac{rl}{u(e)}} \frac{\mathbb{E}_\nu\{g_1\}}{\Delta t} \leq a^{\frac{rl}{u(e)}} e^{-\Delta t \bar{X}} \bar{X} a^{\frac{r}{u(e)}} e^{\Delta t \bar{X} a^{\frac{r}{u(e)}}}$$

which together with the definition of $\bar{X}$ and $\Delta t \leq 1$ proves the lemma. ∎

Lemma 3.8 *For every constant b*

$$\mathbb{E}\{lb^l\} < \infty$$

$$\mathbb{E}\{b^l\} < \infty$$

Proof:

$$\begin{aligned}
\mathbb{E}\{lb^l\} &= \sum_{j=0}^{\infty} jb^j \mathrm{P}rob(l = j) = \sum_{j=0}^{\infty} jb^j e^{-\bar{l}} \frac{\bar{l}^j}{j!} \\
&= e^{-\bar{l}} b\bar{l} \sum_{j=1}^{\infty} \frac{(b\bar{l})^{j-1}}{(j-1)!} = e^{-\bar{l}} b\bar{l} e^{b\bar{l}} < \infty
\end{aligned}$$

The proof for the second case follows using the same calculations. ∎

Consider the case where at time 0 the system is in some fixed state f_0, and at time t the system is in some (random) state f_t, with an associated probability measure μ. We assume that the transition from f_0 to f_t occured due to random events (arrivals and terminations) in the interval $[0, t]$ together with routing decisions made by our routing algorithm. As before, let f^* be any flow that satisfies the demands $d_i = \lambda_i r_i T$.

Lemma 3.9 *Assume that*

$$\Psi(f_0) \leq \Psi\left((1 + 3\epsilon)f^*\right).$$

Then, for every $t > 0$, we have

$$\mathbb{E}\{\Psi(f_t)\} \leq \Psi\left((1 + 3\epsilon)f^*\right).$$

Proof: If the claim of the theorem does not hold, then at some time t we have $\mathbb{E}\{\Psi(f_t)\} > \Psi\left((1 + 3\epsilon)f^*\right)$. Consider maximum $t_0 < t$ such that $\mathbb{E}\{\Psi(f_{t_0})\} = \Psi\left((1 + 3\epsilon)f^*\right)$. The existence of such t_0 follows from the continuity of the function, which is implied by the differentiability. Then

$$0 < \mathbb{E}\{\Psi(f_t)\} - \mathbb{E}\{\Psi(f_{t_0})\} = \int_{s=t_0}^{t} \frac{\partial}{\partial t} \mathbb{E}\{\Psi(f_t)\}\Big|_{t=s} ds$$

By construction, for $s \in [t_0, t]$, we have $\mathbb{E}\{\Psi(f_s)\} > \Psi\left((1 + 3\epsilon)f^*\right)$. Theorem 3.6 implies that for $s \in [t_0, t]$ the derivative $\frac{\partial}{\partial t}\mathbb{E}\{\Psi(f_t)\}\big|_{t=s} < 0$, leading to a contradiction. ∎

Theorem 3.10 *Assume that for every edge e, we start in a state where* $\Lambda_e(f_0) \leq \Lambda_e(f^*)$. *If the granularity condition (2) is satisfied, then the routing algorithm maintains that for every* $t > 0$, $\Lambda(f_t) \leq \Lambda(f^*) + 4\epsilon$ *with probability greater than* $1 - p$.

Proof: Markov's inequality, together with Lemma 3.9, implies

$$\mathrm{P}rob\left\{\Psi(f_t) \leq p^{-1}\Psi\left((1+3\epsilon)f^*\right)\right\} \geq 1 - p.$$

Recall that we have chosen $a = (p^{-1}m)^{1/\epsilon}$. Using the assumption that $\Lambda(f^*) \leq 1$, together with Lemma 3.2, implies that, with probability of at least $1 - p$,

$$\Lambda(f_t) \leq \frac{\log p^{-1}m}{\log a} + \Lambda((1+3\epsilon)f^*) \leq 4\epsilon + \Lambda(f^*).$$

▮

4 Throughput and Profit Maximization

In this section we address the throughput-maximization model. Each circuit is associated with a profit parameter ρ. The profit accrued as a result of routing a circuit is $\rho r t$, where r is the bandwidth and t is the actual holding time of the circuit. We assume that there exist k types of circuits, where a circuit of type i arrives with rate λ_i, requires a path of bandwidth r_i between v_i and w_i, and is associated with a profit parameter ρ_i. As before, we assume that the arrival rates λ_i are unknown.

A routing and admission control algorithm should decide whether to accept or reject a circuit; if the circuit is accepted, enough bandwidth should be reserved along a path between the source and destination for the duration of the circuit. The total bandwidth reserved on every edge should not exceed its capacity. (This is in contrast to the congestion model discussed in the previous section, where we allow to exceed capacity but do not allow rejections.) The performance measurement is the average profit generated by the accepted circuits, where the average is taken with respect to the distribution of the input events.

The profit-maximization algorithm is shown in Figure 1. We will assume that the algorithm knows the average rate of profit generated by all incoming circuits $D = \sum \lambda_i r_i \rho_i T$ and the average duration T. We will also assume that the algorithm is given some target feasible rejection rate R^* (i.e. R^* is the expected fraction of the profit lost by an offline algorithm).

We translate the problem into a congestion-minimization problem by adding an additional "edge" rej. Each circuit is either routed on a path between its endpoints, or is assigned to rej. The capacity of rej is given by $u(rej) = (R^* + O(\epsilon))D$. The algorithm keeps track of all the circuits, both rejected and accepted. Since termination times of rejected circuits never become known, the algorithm picks them using the same distribution as the termination times of the accepted circuits.

The current cost of rej is computed using the same formula as the costs of the original edges:

$$cost(rej) = \frac{1}{u(rej)} a^{\Lambda_t(rej)}$$

where the congestion at rej is defined by

$$\Lambda_t(rej) = \frac{1}{u(rej)} \sum_i \rho_i f_t^i(rej).$$

The decision whether to assign to rej or to a shortest-path route depends on their relative costs. Circuits assigned to rej are rejected. Circuits assigned to a path with insufficient available capacity

Given a new circuit (v, w, r, ρ) compute:

cost-accept $\leftarrow$ Cost of (v, w) shortest path
w.r.t. the cost $a^{\Lambda_t(e)}/u(e)$.

cost-reject $\leftarrow \rho a^{\Lambda_t(rej)}/u(rej)$

If **cost-reject** $<$ **cost-accept then**
Reject the circuit.
Add the circuit to $f(rej)$.
Compute a random termination time.

Else (**cost-reject** $>$ **cost-accept**)
Update f as if the circuit is routed along the shortest path.
If (routing will cause capacity violation)
Reject the circuit.
Compute a random termination time.
Else
Route the circuit

Figure 1: Profit maximization algorithm

are rejected as well. The rest of the circuits are accepted. Note that f_t^i counts both the rejected and the accepted circuits.

Observe that, as in Theorem 3.1, R^*D is bounded by solution to the weighted max-sum multicommodity flow problem with demands $\lambda_i r_i T$, where flow corresponding to type i requests is weighted by ρ_i.

In order to address the profit-maximization model, we will use the fact that the algorithm presented in the previous section can be used to solve a more general problem. More precisely, the capacity β_i^e reserved for a circuit of type i on edge e can be made to depend on the edge and circuit-type pair.

Theorem 4.1 *The generalized congestion-minimization algorithm has the same performance as stated in Theorem 3.10 if the following generalized granularity condition is satisfied.*

$$\frac{\beta_i^e}{u(e)} \leq \frac{\epsilon^2}{\log p^{-1} m}$$

The proof of the theorem is identical to the original proof. Note that setting $\beta_i^e = r_i$ corresponds to the algorithm described in the previous section.

The new granularity conditions needed for profit-maximization algorithm are as follows:

$$(4) \quad \begin{aligned} \frac{r_i}{u(e)} &\leq \frac{\epsilon^2}{\log \epsilon^{-1}(m+1)} \\ \frac{\rho_i r_i}{D} &\leq \frac{\epsilon^2}{\log \epsilon^{-1}(m+1)} \end{aligned}$$

For simplicity, we assume that at time 0, the system is empty. (This assumption can be relaxed as for congestion model.)

Theorem 4.2 *If the granularity conditions (4) hold, and $a = (\epsilon^{-1}m)^{1/\epsilon}$, the expected fraction of rejected profit achieved by the profit maximization algorithm is bounded by $R^* + O(\epsilon)$, assuming $\epsilon \leq 1/3$.*

Proof: Set $u(rej) = (R^* + 18\epsilon)D$. Better constants can be achieved if we assume that $\epsilon \ll 1/3$.

Note that rejections are caused by two different conditions: assignment of a circuit to rej or assignment to an over-congested path. We will estimate the average loss of profit in each one of these cases separately.

Let f^* be a solution to a weighted max-sum multicommodity flow problem with demands $d_i = \lambda_i r_i T$. The flow f^* satisfies demands $d'_i \leq d_i$, where $\sum_i (d_i - d'_i)\rho_i = R^* D$. Define f^{**} on the graph together with the additional edge rej as follows:

$$\begin{aligned} {f^{**}}^i(e) &= (1-4\epsilon){f^*}^i(e) \\ {f^{**}}^i(rej) &= d_i - (1-4\epsilon)d'_i \end{aligned}$$

Observe that f^{**} satisfies capacities $(1-4\epsilon)u(e)$ and satisfies demands $d'_i(1-4\epsilon)$. Thus, by Lemma 3.9, we have for every t:

$$\mathbb{E}\{\Psi(f_t)\} \leq \Psi((1+3\epsilon)(f^{**}))$$

Markov inequality implies that, with with probability $1-\epsilon$,

$$\Psi(f_t) \leq \epsilon^{-1}\Psi((1+3\epsilon)(f^{**}))$$

Lemma 3.2 implies that with probability $1-\epsilon$:

$$\begin{aligned} \Lambda_t(f_t) &\leq \frac{\log \epsilon^{-1}(m+1)}{\log a} + (1+3\epsilon) \max_{e \in \{E, rej\}} \Lambda(f^{**}(e)) \\ &\leq \epsilon + (1+3\epsilon) \max \left[1-4\epsilon, \frac{(R^*+4\epsilon)D}{(R^*+18\epsilon)D}\right] \\ &\leq \epsilon + (1+3\epsilon) \max \left[1-4\epsilon, \frac{1+4\epsilon}{1+18\epsilon}\right] \\ &\leq \epsilon + 1 - \epsilon = 1 \end{aligned}$$

Thus, the expected lost profit due to assignment of circuits to rej is bounded by $(1-\epsilon)u(rej) + \epsilon D \leq (R^* + 19\epsilon)D$. The average lost profit rejected due to assignment to over-congested edges is bounded by ϵD because the probability for the system to be congested at a given time is bounded by ϵ, and the input is independent of the current state. The claim follows. ∎

5 Extensions

5.1 Changes in the Traffic Matrix

The model described in this paper assumes a fixed traffic matrix. In this case, a possible alternative to our algorithm is to learn (estimate) this traffic matrix and then route according to a fractional multicommodity flow solution. In practice it is possible that the traffic matrix will change too fast for this approach to be practical. In contrast to methods that relay on gathering of statistical information, our algorithm has the same performance even if the traffic matrix changes in discrete intervals. For example, in the congestion minimization case, our algorithm achieves congestion that is within $O(\epsilon)$ of the maximal congestion over all traffic matrices used to generate the input.

To observe that the same estimates works when the traffic matrix is changing in some discrete intervals we should note the following two facts:

- In Corollary 3.4 it is enough to assume that the traffic matrix is fixed for some small interval $[t_0, t]$.
- In Lemma 3.9 it is enough to assume that $\mathbb{E}\{\Psi(f_t)\}$ is a piecewise differential function.

5.2 Computation with Rounding Errors

Finding the exact costs on each edge might involve an unreasonable number of bit operations. In this subsection we will show that working with some finite precision in a floating point model has limited effect on the performance of the algorithm.

Theorem 5.1 *If computations of costs is done up to some $1 \pm \delta$ precision then all the results apply, with an addition factor of $(1+\delta)^2$.*

Proof: Let $\bar{c}(e) = a^{\Lambda(e)}/u(e)$ and $\tilde{c}$ be some finite precision approximation to $\bar{c}$ such that $1 - \delta < \bar{c}(e)/\tilde{c}(e) < 1 + \delta$. Let $\bar{f}$ be as in theorem 3.5, and $\tilde{f}$ the equivalent flow which is the minimal cost flow with respect to $\tilde{c}$. Since the algorithm routes flow using $\tilde{f}$ rather then $\bar{f}$, following the proof of theorem 3.5 we know that

$$\frac{\partial}{\partial t}\,\mathbb{E}\{\Psi(f_t)\}|_{t=t_0} \le \log a \frac{1}{T} \sum_e a^{\Lambda_{t_0}(e)} \left[\left(\Lambda(T\tilde{f}(e)) - \Lambda_{t_0}(e)\right) + \eta \left(\Lambda(T\tilde{f}(e)) + \Lambda_{t_0}(e)\right)\right]$$

Using the definitions of $\bar{c}, \tilde{c}, \bar{f}, \tilde{f}$ we know that

$$\sum_e \bar{c}(e)\tilde{f}(e) \le (1+\delta)\sum_e \tilde{c}(e)\tilde{f}(e) \le (1+\delta)\sum_e \tilde{c}(e)\bar{f}(e) \le (1+\delta)^2 \sum_e \bar{c}(e)\bar{f}(e)$$

Using the rest of the proof of 3.5 we get that

$$\frac{\partial}{\partial t}\,\mathbb{E}\{\Psi(f_t)\}|_{t=t_0} \le \frac{1}{T}(1-\eta)\left[\Psi\left((1+\delta)^2(1+3\epsilon)f^*\right) - \Psi(f_{t_0})\right]$$

All the other theorems apply without any change. ▮

Acknowledgements

We would like to thank Debasis Mitra, Ram Ramakrishnan, and Éva Tardos for many helpful discussions.

References

[1] G.R. Ash. Use of trunk status map for real-time DNHR. In *Proc. of ITC-12, Torino, Italy*, 1988.

[2] G.R. Ash, R.H. Cardwell, and R.P. Murray. Design and optimization of networks with dynamic routing. *Bell Syst. Tech. J.*, 60:1787–1820, 1981.

[3] G.R. Ash, J.S. Chen, A.E. Frey, and B.D. Huang. Real-time network routing in an integrated network. In *Proc. of ITC-13, Copenhagen, Denmark*, 1991.

[4] G.R. Ash, A.H. Kafker, and K.R. Krishnana. Servicing and real-time control of networks with dynamic routing. *Bell Syst. Tech. J.*, 60:1821–1845, 1981.

[5] J. Aspnes, Y. Azar, A. Fiat, S. Plotkin, and O. Waarts. On-line machine scheduling with applications to load balancing and virtual circuit routing. In *Proc. 25th Annual ACM Symposium on Theory of Computing*, pages 623–631, May 1993.

[6] B. Awerbuch, Y. Azar, and S. Plotkin. Throughput competitive on-line routing. In *Proc. 34th IEEE Annual Symposium on Foundations of Computer Science*, pages 32–40, November 1993.

[7] B. Awerbuch, Y. Azar, S. Plotkin, and O. Waarts. Competitive routing of virtual circuits with unknown duration. In *Proc. 5th ACM-SIAM Symposium on Discrete Algorithms*, pages 321–327, January 1994.

[8] B. Awerbuch, R. Gawlick, T. Leighton, and Y. Rabani. On-line admission control and circuit routing for high-performance computing and communication. In *Proc. 35th IEEE Annual Symposium on Foundations of Computer Science*, pages 412–423, November 1994.

[9] Y. Azar, A. Broder, and A. Karlin. On-line load balancing. In *Proc. 33rd IEEE Annual Symposium on Foundations of Computer Science*, pages 218–225, 1992.

[10] Y. Azar, B. Kalyanasundaram, S. Plotkin, K. Pruhs, and O. Waarts. On-line load balancing of temporary tasks. In *Proc. Workshop on Algorithms and Data Structures*, pages 119–130, August 1993.

[11] A. Borodin, N. Linial, and M. Saks. An optimal online algorithm for metrical task systems. *J. ACM*, (39):745–763, 1992.

[12] J. Garay, I. Gopal, S. Kutten, Y. Mansour, and M. Yung. Efficient on-line call control algorithms. In *Proc. of 2nd Annual Israel Conference on Theory of Computing and Systems*, 1993.

[13] J.A. Garay and I.S. Gopal. Call preemption in communication networks. In *Proc. INFOCOM '92*, volume 44, pages 1043–1050, Florence, Italy, 1992.

[14] R. Gawlick, A. Kamath, S. Plotkin, and K. Ramakrishnan. Routing of virtual circuits with unknown holding times. Unpublished manuscript, 1994.

[15] R. Gawlick, A. Kamath, S. Plotkin, and K. Ramakrishnan. Routing and admission control in general topology networks. Technical Report STAN-CS-TR-95-1548, Stanford University, 1995.

[16] R.J. Gibbens and F. P. Kelly. Dynamic routing in fully connected networks. *IMA J. Math. Contr. Inform.*, 7:77–111, 1990.

[17] D. Karger and S. Plotkin. Adding multiple cost constraints to combinatorial optimization problems, with applications to multicommodity flows. In *Proc. 27th Annual ACM Symposium on Theory of Computing*, pages 18–25, May 1995.

[18] A.R. Karlin, M.S. Manasse, L.Rudolph, and D.D. Sleator. Competitive snoopy caching. *Algorithmica*, 1(3):70–119, 1988.

[19] F. P. Kelly. Blocking probabilities in large circuit-switched networks. *Adv. Appl. Prob.*, 18:473–505, 1986.

[20] P. Klein, S. Plotkin, C. Stein, and É. Tardos. Faster approximation algorithms for the unit capacity concurrent flow problem with applications to routing and finding sparse cuts. *SIAM Journal on Computing*, 23(3):466–487, June 1994.

[21] J. Kleinberg and É. Tardos. Disjoint path in densely embedded graphs. In *Proc. 36th IEEE Annual Symposium on Foundations of Computer Science*, 1995.

[22] K.R. Krishnan and T.J. Ott. Forward-looking routing: A new state-dependent routing scheme. In *Proc. of ITC-12, Torino, Italy*, 1988.

[23] T. Leighton, F. Makedon, S. Plotkin, C. Stein, É. Tardos, and S. Tragoudas. Fast approximation algorithms for multicommodity flow problems. *J. Comp. and Syst. Sci.*, 50:228–243, 1995. (Invited paper).

[24] Y. Ma and S. Plotkin. Improved lower bounds for load balancing of tasks with unknown duration. Unpublished manuscript, 1995.

[25] M.S. Manasse, L.A. McGeoch, and D.D. Sleator. Competitive algorithms for online problems. In *Proc. 20th Annual ACM Symposium on Theory of Computing*, pages 322–332, 1988.

[26] V. Marbukh. An asymptotic study of a large fully connected communication network with reroutes. *Problemy Peredachi Informatsii*, 3:89–95, 1981.

[27] V. Marbukh. Study of a fully connected communication network with many nodes and circumventing routes. *Automatica i Telemechanica*, 12:86–94, 1983.

[28] D. Mitra, R.J. Gibbens, and B.D. Huang. State-dependent routing on symmetric loss networks with trunk reservations – I. *Annals of Operations Research*, (35):3–30, 1992.

[29] D. Mitra, R.J. Gibbens, and B.D. Huang. State-dependent routing on symmetric loss networks with trunk reservations – II. *IEEE Trans. on Communications*, 41(2), February 1993.

[30] T.J. Ott and K.R. Krishnan. State-dependent routing of telephone traffic and the use of separable routing schemes. In *Proc. of ITC-11, Kyoto, Japan*, 1985.

[31] S. Plotkin. Competitive routing in ATM networks. *IEEE J. Selected Areas in Comm.*, pages 1128–1136, August 1995. Special issue on Advances in the Fundamentals of Networking. (Invited paper).

[32] S. Plotkin, D. Shmoys, and É. Tardos. Fast approximation algorithms for fractional packing and covering problems. In *Proc. 32nd IEEE Annual Symposium on Foundations of Computer Science*, pages 495–504, October 1991.

[33] F. Shahrokhi and D. Matula. The maximum concurrent flow problem. *J. Assoc. Comput. Mach.*, 37:318–334, 1990.

[34] S. Sibal and A. DeSimone. Controlling alternate routing in general-mesh packet flow networks. Technical Report BL045370F-940603-01TM, AT&T, June 1994.

[35] D.D. Sleator and R.E. Tarjan. Amortized efficiency of list update and paging rules. *Comm. ACM*, 28(2):202–208, 1985.

[36] N. Young. Randomized rounding without solving the linear program. In *Proc. 6th ACM-SIAM Symposium on Discrete Algorithms*, pages 170–178, 1995.