# An Approximate Calculation of Max-Min Fair Throughputs for Non-Persistent Elastic Flows

Pinaki Shankar Chanda, Anurag Kumar, Arzad Alam Kherani
Dept. of Electrical Communication Engg.
Indian Institute of Science, Bangalore, 560012, INDIA

***Abstract*—The general problem we consider is the analysis of a model in which there are several routes in a network, on each route elastic flows arrive randomly according to some arrival process, and each flow transfers a finite volume of data sampled from some distribution. We are interested in computing a measure of average flow throughput on each route, for a given bandwidth sharing mechanism. Such models arise in problems of network dimensioning and traffic engineering.**

**In this paper, we assume Poisson arrivals of file transfer requests on each route, the transfer volumes are fluid and arbitrarily distributed. At each instant the network shares the bandwidth among the ongoing flows according to the max-min fair bandwidth sharing mechanism, i.e.,** *Instantaneous Max-Min Fair* **(IMMF) sharing. The measure of performance we consider is the time average bandwidth obtained by flows on each route. We propose a heuristic algorithm for obtaining an approximation for this performance measure for arbitrary routes in an arbitrary network topology. Simulations with various network topologies are used to evaluate the proposal. In spite of its simplicity, we find that the approximation works quite well in a variety of topologies that we have studied.**

## I. Introduction

Traffic engineering of networks requires effective models that can predict the performance seen by end users. An important performance measure is the average throughput of flows belonging to a particular route. Most studies related to throughput analysis ([1], [2], [3]) have assumed *persistent* flows, i.e., the flows transfer an infinite amount of data. In reality, new sessions start, transmit some finite number of packets and then close their connections. Hence the consideration of a flow arrival process and random transfer volumes in modeling elastic flows in a network is important. In this situation, the throughputs of the flows in a route depend on the random number of ongoing transfers on the various routes, and on the network topology.

In the work presented here, we consider an arbitrary network topology with fixed routes. On each route, we assume that transfer requests, or flows, arrive in a Poisson process. Each request is for the transfer of a random volume of data, and the sequence of transfer volumes are independent and identically distributed. We assume that *instantaneous* max-min fair (IMMF) sharing is achieved in the network; i.e., at each instant of time the rate obtained by a flow on a route is its max-min fair share given the ongoing flows at that time. Such models have been proposed in [4], [5], and [6]. We adopt the average bandwidth share of the flows on different routes in the network as our performance index. In this modeling framework, our contribution is an algorithm to approximately calculate the average bandwidth share of flows on each route. Simulations for different network topologies have been performed to evaluate this proposal.

In ATM networks, elastic flows are carried on the ABR (available bit rate) service, and the bandwidth sharing objective is to achieve max-min fairness. If the bandwidth-delay product is small compared to the flow durations then the ABR control loops would converge relatively quickly compared to the time scales of variation of the number of flows on the routes, and approximately IMMF sharing would hold (see, for example, [7]).

In an internet, bandwidth sharing is achieved by TCP's end-to-end adaptive window mechanism. Recent research has suggested that TCP's AIMD rate control, based on packet-loss feedback, is equivalent to a bandwidth sharing objective called $F_A^h$ fairness [8]. It has been shown in [9] that TCP Vegas achieves proportional fairness; see also [10]. In some simple situations (simple topologies or a single bottleneck for each route) proportional fairness is equivalent to max-min fairness. Further, TCP can be viewed as a whole class of adaptive window based algorithms, and the bandwidth sharing that is achieved depends on the network feedback and source adaptation. With this point of view, it can be seen that with appropriate feedback from routers TCP can actually achieve max-min fairness ([11]). Each link computes its MMF parameter $\eta_l,\ l \in L$ (the set of all links) (see [7]), and marks packets passing through it with probability $e^{-\alpha\eta_l}$, for some $\alpha > 0$. For large $\alpha$ it follows that the total marking probability along route $r$ is approximately $e^{-\alpha \min_{l\in L_r} \eta_l}$, where $L_r$ is the set of all links carrying route $r$ (see also [12]). Thus each source on route $r$ effectively learns $\min_{l\in L_r} \eta_l$ (its fair rate), measures the round trip time along the path and adjusts its window accordingly (see also [13]).

Some prior literature related to our work is the following. The importance of traffic engineering models with non-persistent flows has been highlighted in [14]. In [4] the authors consider the same stochastic model as we do, and for exponentially distributed transfer volumes they provide an ergodicity condition for the process of the number of flows on each route. The fairness objectives considered are: max-min fairness, weighted max-min fairness and proportional fairness. These results are generalised in [5] to a broad class of fairness

0-7803-7206-9/01/$17.00 © 2001 IEEE

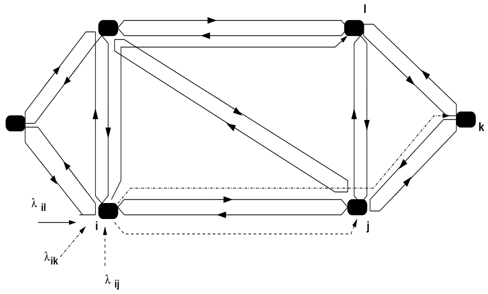


Fig. 1. A backbone network showing bidirectional links, and several routes on which finite volume flows arrive randomly; also shown are routes (i,j), (i,k) and (i,l).

objectives, of which max-min fairness is a special case. Computational results are provided for some simple topologies. In [6] a bandwidth sharing objective called "min bandwidth sharing" is introduced, and a technique is studied for performance analysis with this policy in large networks with certain structure. The papers [4], [5], and [6] all make the assumption that the bandwidth shares adjust instantaneously as flows arrive and depart. This can be expected to be a good approximation to the actual network behaviour when the delay-bandwidth product is small relative to the transfer volumes.

There are a few papers on performance models of TCP controlled bandwidth sharing, with finite volume flows. Simple single link networks are studied via product form queueing models in [15] and [16]. A detailed model for an additive-increase and multiplicative-decrease congestion control protocol, with session arrivals and departures is provided in [17]. An iterative technique for networks is proposed in [18].

The difficulty of exact computation of performance in our model has been recognised in [19] and [20]. In the latter thesis some approximations are provided for simple topologies.

## II. The Model and Notation

**Network and Traffic Model:** We consider a backbone packet network of a service provider or an enterprise, and assume that this backbone network is the bottleneck for the traffic it carries. There are several fixed, unidirectional routes between the edge routers. The source router of each of these routes sees a random arrival process of initiations of the transfer of random volumes of data to the destination router of the route. We depict this scenario in Figure 1. For a route with source node $i$ and destination node $j$, $\lambda_{i,j}$ is the rate of arrival of flows from $i$ to $j$. Each bidirectional link of the network is represented as a pair of separate unidirectional links carrying traffic in opposite directions, i.e., we think of the network as a directed graph as shown in Figure 1. In this work when we use the term "link" we mean a unidirectional link.

**Modeling assumptions:** Flow arrival instants on each route constitute a Poisson process; the transfer volumes are independent and identically distributed and are independent of the arrival process. The Poisson approximation for the arrival process of file transfer connection initiations has been found to be reasonable in [21]. It may also be reasonable to model the initiations of web sessions as a Poisson process, considering that the backbone sees arrivals from a large number of independent users. Further, the Poisson flow initiation model is mathematically tractable and has been used by several researchers ([4], [5]). We assume that the network bandwidth is shared among the flows at time $t$ in the IMMF fashion, as defined in the introduction. Note that for the case of a single link, IMMF reduces to bandwidth sharing in the Processor Sharing manner.

Theorem 3.1 of [4] and Theorem 1 of [5] state that the vector random process of the number of flows on the various routes converges in distribution if the total offered bit rate on any link is less than the link capacity; in [5] it is also asserted that this result holds for arbitrary transfer volume distribution.

**Notation:** In a network with $R$ fixed routes and a fixed set $\mathcal{L}$ of $L$ links, we use the following notation.
$\mathcal{R}$: set of all routes in the network; $|\mathcal{R}| = R$
$\lambda_r$: rate of flow arrivals on route $r \in \mathcal{R}$
$V_r$: random variable denoting the transfer volumes on route $r$, $r \in \mathcal{R}$; it has mean $E(V_r)$
$C_j$: the capacity of link $j \in L$
$\rho_r^j$: normalised load on link $j$ from route $r$ that uses link $j$; i.e., $\rho_r^j = \frac{\lambda_r E(V_r)}{C_j}$
$N_r(t); r \in \mathcal{R}$: the number of flows in route $r$ at time $t$
$(N_1, N_2 ... N_R)$: stationary random vector for $\{N_r(t) : r \in \mathcal{R}\}$
$\pi(n_1, n_2 ... n_R)$: stationary distribution of $(N_1, N_2 ... N_R)$
In the case of a single link network with a single route (or single "class"), we use this notation without the subscripts and superscripts. The notation also applies to a single link network with several routes (or classes).

## III. The Performance Measure

### A. Single Class Traffic on a Link

On a single link, of bit rate $C$, IMMF bandwidth sharing reduces to an M/G/1 processor sharing model. There are various measures of flow throughput (refer [17] for a detailed description). We consider the time averaged bandwidth share per ongoing flow, conditioned on there being at least one flow, as our performance measure. This was introduced in [20] and is defined as,

$$\sigma := \lim_{t \to \infty} \frac{\int_0^t \frac{C}{N(u)} I_{\{N(u)>0\}} du}{\int_0^t I_{\{N(u)>0\}} du} \quad (1)$$

where $I_{\{N(u)>0\}} = \begin{cases} 1 & \text{if } N(u) > 0 \\ 0 & \text{otherwise} \end{cases}$. On a single link, $\sigma$ has been found to approximate the average flow throughput well while being computationally tractable [17]. For $\rho < 1$ it is easy to see that,

$$\sigma = \frac{\sum_{n=1}^{\infty} C \frac{\pi(n)}{n}}{\rho} \quad (2)$$

For the M/G/1 PS model $\pi(n) = (1-\rho)\rho^n$ and hence,

$$\frac{\sigma}{C} = \frac{1-\rho}{\rho} \ln \frac{1}{1-\rho} \tag{3}$$

Since $\pi$ is insensitive to the distribution of $V$, so is $\sigma$.

### B. Multi-Class Traffic on a Link

When multiple routes share a link we view the link as being shared by multiple traffic classes. Define, $n = n_1 + n_2 + ...... + n_R$. The steady state distribution of $(N_1, N_2 ....N_R)$ is given by [22]

$$\pi(n_1, n_2 ....n_R) = (1-\rho)\, n! \prod_{r=1}^{R} \frac{\rho_r^{n_r}}{n_r!} \tag{4}$$

where $\rho = \sum_{i=1}^{R} \rho_i$. The time average per flow bandwidth share of class $r$, i.e., $\sigma_r, r \in \mathcal{R}$, is defined as the time averaged rate per ongoing flow, conditioned on there being at least one class $r$ flow. Hence,

$$\sigma_r := \lim_{t\to\infty} \frac{\int_0^t \frac{C}{N(u)} I_{\{N_r(u)>0\}} du}{\int_0^t I_{\{N_r(u)>0\}} du} \tag{5}$$

where $N(u) = \sum_{r\in\mathcal{R}} N_r(u)$, i.e., the total number of flows at time $u$ on this link. Note the difference between $\sigma_r$ and $\sigma$ as defined earlier for single class traffic. For $\rho < 1$,

$$\sigma_r = \frac{\sum_{(n_1,n_2,....n_R):n_r>0} \frac{C}{(n_1+n_2+....+n_R)} \pi(n_1, n_2 ....n_R)}{\sum_{(n_1,n_2,....n_R):n_r>0} \pi(n_1, n_2, ..., n_R)}$$

and we can easily simplify this to (details are given in [23]),

$$\sigma_r = C \frac{(1-\rho)(1-\bar{\rho_r})}{\rho_r} \ln(\frac{1-\bar{\rho_r}}{1-\rho}) \tag{6}$$

where $\bar{\rho}_r = \sum_{\forall i\in\mathcal{R}:i\neq r} \rho_i$. Note that when there is only one class of traffic, $\bar{\rho_1} = 0$, $\rho_1 = \rho$ and Equation 6 yields $\sigma_1 = C\frac{1-\rho}{\rho} \ln(\frac{1}{1-\rho})$, which is the same as $\sigma$, as expected.

## IV. Approximately Calculating $\sigma_r$ in a Network

We define the *per flow bandwidth share of a route* in a network as the time averaged IMMF rate per flow belonging to a particular route, conditioned on there being at least one flow active on that route. Denoting this by $\sigma_r$, we see that

$$\sigma_r := \lim_{t\to\infty} \frac{\int_0^t x_r(u) I_{\{N_r(u)>0\}} du}{\int_0^t I_{\{N_r(u)>0\}} du} \tag{7}$$

*where $x_r(u)$ is the max-min fair rate of a flow in route $r$ in the network at time $u$*; this, of course, depends on the vector $(N_1(u), N_2(u), \cdots, N_r(u))$. To get a closed form expression for $\sigma_r$ in a network we need to know the stationary distribution of the number of flows in different routes. This is a difficult problem to solve ([6], [19], [20]). We, therefore, propose a heuristic algorithm to obtain $\sigma_r, r \in \mathcal{R}$.

### A. Motivation: Obtaining $\sigma_r$ Recursively on a Single Link

Consider again a single link, and let, $\sigma_k(C; \rho_1, \rho_2, ....\rho_R)$ denote $\sigma_k$ with total link capacity $C$ and normalized load $\rho_j$ on route $j$, $1 \le j \le R$. As shown above,

$$\begin{aligned} &\sigma_k(C; \rho_1, \rho_2, ....\rho_R) \\ &= C\frac{(1-\rho)(1-\bar{\rho_k})}{\rho_k} \ln(\frac{1-\bar{\rho_k}}{1-\rho}) \end{aligned} \tag{8}$$

with $\bar{\rho_k} = \sum_{i=1:i\neq k}^{R} \rho_i$.

*Theorem IV.1:* Consider a given indexing of the $R$ routes. Then for $k \in \{1, 2...R\}$
$\sigma_k(C; \rho_1, \rho_2, ....\rho_R) = \sigma_k(C(1 - \sum_{i=1}^{k-1} \rho_i); \rho_k, \rho_{k+1}...\rho_R)$
□

The proof of this theorem involves trivial checking (see [23]). The above theorem implies the following recursive approach for calculating $\sigma_k, 1 \le k \le R$, on a single link, paralleling the approach for persistent flows (as in [24]).

**Recursive algorithm on a single link:** Fix an indexing of the routes $k, 1 \le k \le R$. Then compute $\sigma_1 = \sigma_1(C; \rho_1, \rho_2, ....\rho_R)$. The average bandwidth share of the flows in route 1 is given by $\sigma_1$. The total bandwidth utilised by route 1 is $C\rho_1$ (unlike persistent flows this does not depend on the bandwidth shares of the other routes). Using the above theorem, we can then compute $\sigma_2 = \sigma_2(C(1-\rho_1); \rho_2, ....\rho_R)$. *Thus at each step the link capacity is reduced by the amount taken up by the routes examined until that step and the bandwidth shares of flows in the remaining routes are computed assuming that the remaining routes share the reduced link capacity.*

### B. Algorithm for Approximate Calculation of $\sigma_r$ in a Network

In this recursive algorithm we determine average bandwidth share of a single route in each iteration. We remove this route from the topology and adjust the capacity of the links carrying this route. Hence on every iteration the route-link incidence matrix (defined below) and the remaining capacity of the links changes.

#### B.1 Additional Notation Used

$R^{(k)}$*:* total number of remaining routes in the network in iteration $k$. Initially, $R^{(0)} = R$.

$\gamma^{(k)}$*:* the $l$th element of this vector is the remaining capacity of link $l \in L$ of the network (say $\gamma_l^{(k)}$) in iteration $k$; i.e., $\gamma^{(k)} = [\gamma_l^{(k)}]_{(1\times|L|)}$. Initially, $\gamma_l^{(0)} = C_l, \ \forall l \in L$.

$A^{(k)}$*:* the route-link incidence matrix in iteration $k$; i.e., $A^{(k)} = [a_{ij}^{(k)}]_{(|\mathcal{R}|\times|L|)}$ where,

$$a_{ij}^{(k)} = \begin{cases} 1 & \text{if route } i \text{ is crossing link } j \\ 0 & \text{otherwise} \end{cases}$$

$\psi$*:* the $r$th element of this vector is the approximation to the average per flow bandwidth share $\sigma_r$ for route $r \in \mathcal{R}$ in the network

$\rho_i^{j(k)}$: load on link $j$ due to route $i$ normalized to the remaining link capacity in iteration $k$; $\rho_i^{j(k)} = \frac{\lambda_i E(V_i)}{\gamma_j^{(k)}}$, $i \in \mathcal{R}$, $j \in L$ such that $a_{ij}^{(k)} = 1$

$\sigma_i^{j(k)}$: average bandwidth share of route $i$ on link $j$ considering link $j$ in isolation, in iteration $k$

B.2 Algorithm

We take a route $r$ in the network and consider the set of links carrying this route. We calculate $\sigma_r$ considering each link $l$ separately (e.g., for route $r$ on link $l$ this is $\sigma_r^l$). This computation is repeated for all the routes. We now take the minimum of all $\sigma_r^l$, $\forall r \in \mathcal{R}, \forall l \in L$. We assign this minimum value as the average bandwidth share of the corresponding route. Now, this route, say $r_0$, is removed from the routing topology and the capacities of all the links that carry this route is reduced by $\lambda_{r_0} EV_{r_0}$. The above procedure is repeated to find the average bandwidth shares of all the routes. At the $k$th step, $k \geq 1$, the algorithm works as follows

**Step (1):** We consider each link of the topology separately and calculate $\sigma_i^{j(k-1)}$. Thus, in this step, while calculating the average bandwidth shares we consider each link and the routes in it in isolation. Hence, $\sigma_i^{j(k-1)}$, i.e., the average bandwidth shares of flows in route $i$ over link $j$, when link $j$ is considered alone is given as (see Equation 6)

$$\sigma_i^{j(k-1)} = \gamma_j^{(k-1)} \Big(1 - \sum_{r=1:a_{rj}^{(k-1)}=1}^{R} \rho_r^{j(k-1)}\Big) \cdot \frac{\big(1 - \sum_{r=1:r\neq i, a_{rj}^{(k-1)}=1}^{R} \rho_r^{j(k-1)}\big)}{\rho_i^{j(k-1)}} \cdot \ln\left(\frac{1 - \sum_{r=1:r\neq i, a_{rj}^{(k-1)}=1}^{R} \rho_r^{j(k-1)}}{1 - \sum_{r=1:a_{rj}^{(k-1)}=1}^{R} \rho_r^{j(k-1)}}\right) \quad (9)$$

**Step (2):** We find

$$\sigma_{n_{k-1}}^{m_{k-1}} = \min\{\sigma_i^{j(k-1)}, \ \forall i \in \mathcal{R} \ \forall j \in L \ \text{ with } a_{ij}^{(k-1)} = 1\}$$

We say that route $n_{k-1}$ is bottlenecked in link $m_{k-1}$ in some average sense; i.e., the throughput of this route is determined by link $m_{k-1}$. Now we assign this value to $\psi_{n_{k-1}}$ ,i.e., the average bandwidth share of route $n_{k-1}$ is approximated as

$$\psi_{n_{k-1}} = \sigma_{n_{k-1}}^{m_{k-1}}$$

**Step (3):** This route $n_{k-1}$ is now removed from the network. Necessary changes are made in the route-link incidence matrix, and a new route-link incidence matrix $A^{(k)}$ is formed.

$$a_{ij}^{(k)} = \begin{cases} a_{ij}^{(k-1)} & \forall i \in \mathcal{R} \text{ and } \forall j \in L \text{ such that } i \neq n_{k-1} \\ 0 & \text{otherwise} \end{cases}$$

and $R^{(k)} = R^{(k-1)} - 1$.

We reduce the capacity of each link in the network over which the route $n_{k-1}$ flows by an amount of $\lambda_{n_{k-1}} E(V_{n_{k-1}})$. Thus, $\gamma^{(k)} = \{\gamma_j^{(k)}; j \in L\}$ is such that

$$\gamma_j^{(k)} = \begin{cases} \gamma_j^{(k-1)} - \lambda_{n_{k-1}} E(V_{n_{k-1}}) & \forall j \text{ such that } a_{n_{k-1}j}^{(k-1)} = 1 \\ \gamma_j^{(k-1)} & \text{otherwise} \end{cases}$$

**Step (4):** $k$ is incremented, and steps (1) to (4) are repeated until no more routes remain.

## V. Comparison of Analysis with Simulations

### A. Simulation Setup

We have compared the results of our analysis with a simulation of the fluid model whose analysis we are approximating. A network topology with some fixed routes is chosen. We simulate the arrivals of flows on each route, the random transfer volumes, and the ideal IMMF bandwidth sharing. In the simulation we have considered two distributions of transfer volumes: exponential distribution, and hyper-exponential distribution matched to a Pareto distribution. Note that the approximate analysis is insensitive to the distribution of the transfer volumes. This is, of course, known to be exact for a single link.

### B. Example 1: A Line Network

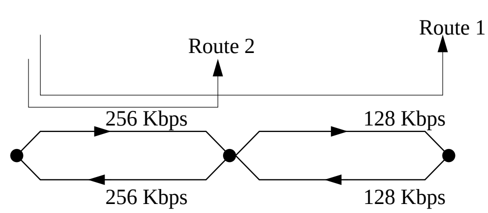


Fig. 2. A two link, low speed network topology with two routes.

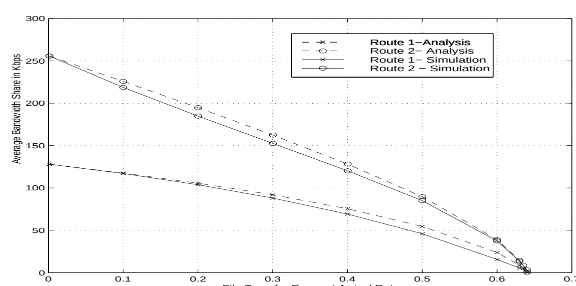


Fig. 3. Average bandwidth share vs. arrival rate for the network shown in Figure 2 ; exponentially distributed flow volumes.

Figure 2 shows a simple low speed network, with two links in tandem, and two routes. The arrival rates of flows to both the routes are taken to be the same. The plot of average bandwidth shares (with expontially distributed flow volumes having mean

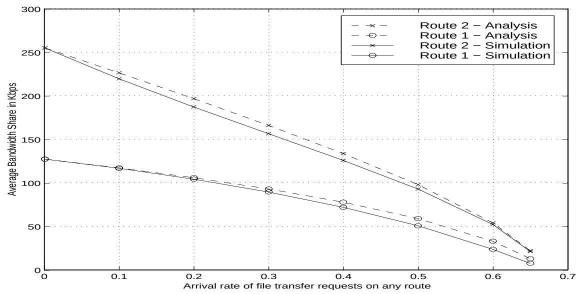


Fig. 4. Average bandwidth share vs. arrival rate for the network shown in Figure 2; flow volumes distributed as a mixture of 7 exponentials matched with a Pareto distribution.

$25$ Kbytes) is shown in Figure 3. Observe that a flow on route $1$ can get a share of at most $128$ Kbps while on route $2$ a flow can get at most $256$ Kbps. Notice that when the arrival rates go to $0$, as expected, the flows on each route do get these maximum possible shares. The bandwidth shares drop as the arrival rates increase. At the arrival rate of $0.64$ flows per second the links saturate and the bandwidth shares become zero.

Figure 4 shows results for a heavy tailed distribution of flow volumes. We take the flow volume as $V = 25\ Z$Kbytes, where $Z$ has the *Pareto distribution* with complementary distribution given by $F^c(z) = (1+bz)^{-a}$, with $a = 2.2$ and $b = .8333$ (i.e., with mean 1). In the simulation we approximate this heavy-tailed distribution with a mixture of 7 exponential distributions (see [25] for this Pareto distribution and this approximation). We get a hyper-exponential distribution with mean $0.9559$ approximating the Pareto distribution; we use the same mean volume ($0.9559 \times 25$Kbytes) in the analysis. Note that since the mean of the flow volumes is less than 25Kbytes, in Figure 4 at the arrival rate of $0.64$ the bandwidth shares are not zero.

We observe from Figures 3 and 4 that the analysis and the simulations match very well, thus supporting our heuristic.

### *C. Example 2: A Mesh Network*

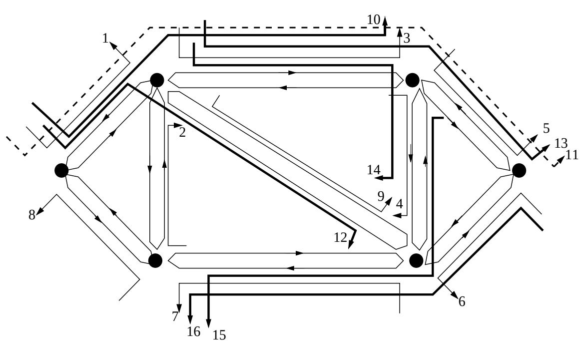


Fig. 5. A mesh network with 9 links and 16 routes. All links are of bandwidth 1Mbps. The numbers in the circles indicate link numbers, the others are route numbers.

We take an arbitrary mesh network with $9$ links and $16$ routes as shown in Figure 5. The simulation and analysis results are plotted in Figure 6 and Figure 7. The arrival rates of flows on all the routes are taken to be the same, and the mean file sizes on all the routes are equal ($25$Kbytes, exponentially distributed). The capacity of all the links is 1Mbps. Here we show the result for a subset of the routes in the network (see [26] for all the simulation results).

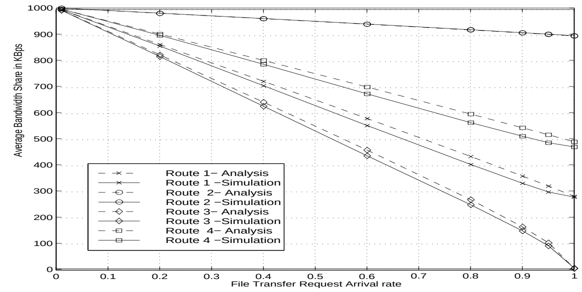


Fig. 6. Average bandwidth shares of flows on routes 1,2,3, and 4, in Kbps vs. arrival rate, for the mesh network in Figure 5; exponentially distributed file sizes

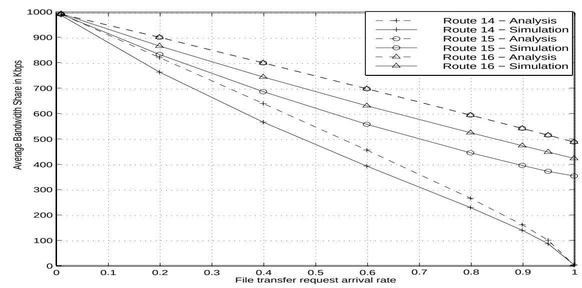


Fig. 7. Average bandwidth shares of flows on routes 14,15, and 16, in Kbps vs. arrival rate for the mesh network shown in Figure 5; exponentially distributed file sizes

Note that the algorithm overestimates the bandwidth share in every case, but is a good approximation. For most of the routes the error is less than $10\%$, with worst case error being $25\%$ (route $15$ in Figure 7).

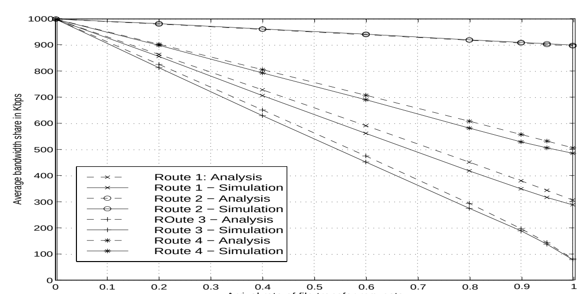


Fig. 8. Average bandwidth share in Kbps of flows on routes 1,2,3, and 4 vs. arrival rate for the mesh network shown in Figure 5; file sizes distributed as a mixture of 7 exponentials matched to a Pareto distribution.

We also show analysis and simulation results using a heavy tailed flow volume distribution (a mixture of 7 exponentials as described above in Section V-B) in Figure 8 and Figure 9.

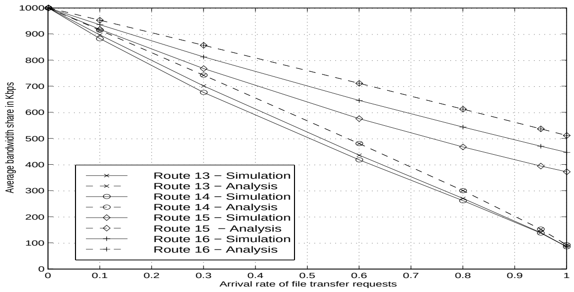


Fig. 9. Average bandwidth share in Kbps of flows on routes 13,14,15, and 16 vs. arrival rate for the mesh network shown in Figure 5; file sizes distributed as a mixture of 7 exponentials matched to a Pareto distribution.

Results are shown for routes 1, 2, 3, 4, 13, 14, 15 and 16. Note that the approximation does as well as for the exponential distribution, and is still an upper bound.

## VI. Conclusion

We have considered a model comprising a network topology with fixed routes, on which elastic flows arrive randomly and transfer random volumes of (fluid) data, while sharing the network in an instantaneous max-min fair fashion. With certain statistical assumptions we have provided a heuristic technique for computing the time average bandwidth shares obtained by the flows on each route. Simulations show that our analysis provides a good approximation. In our current work we are addressing several questions that flow from this paper:

1. Can such a model be directly useful for modelling elastic flows in internets, with certain restrictions, such as very small bandwidth delay products, and high speeds?
2. Does the analysis provably provide an upper bound to $\sigma_r$ for every route $r \in \mathcal{R}$?
3. Can the analysis of a single link carrying several classes of TCP controlled traffic be extended in a similar way (as in this paper) to model an entire network?

## VII. Acknowledgement

This work was supported by a research grant from HFCL (Himachal Futuristic Communications Limited).

## References

[1] Jitendra Padhye, Victor Firoiu, Don Towsley, and Jim Kurose, "Modeling TCP Throughput: A Simple Model and its Empirical Validation," *SIGCOMM'98*, 1998.

[2] Anurag Kumar, "Comparative Performance Analysis of Versions of TCP in Local Network with a Lossy Lnk," *IEEE/ACM Transactions on Networking*, vol. 6, no. 4, August 1998.

[3] T.V. Lakshman and Upamanyu Madhow, "The Performance of TCP/IP for Networks with High Bandwidth Delay Products and Random Loss," *IEEE/ACM Transactions on Networking*, vol. 5, no. 3, pp. 336–350, June 1997.

[4] Gustavo De Veciana, Tae-Jin Lee, and Takis Konstantopoulos, "Stability and Performance Analysis of Network Supporting Elastic Services," *IEEE/ACM Transactions on Networking*, vol. 9, no. 1, pp. 2–14, February 2000.

[5] Thomas Bonald and Laurent Massoulie, "Impact of Fairness on Internet Performance," in *Proc. ACM Sigmetrics*. ACM, 2001.

[6] Guy Fayolle, Arnaud de La Fortelle, Jean-Marc Lasgouttes, Luarent Massoulie, and James Roberts, "Best-Effort Networks: Modeling and Performance Analysis via Large Networks Asymptotics," in *Proc. IEEE Infocom*. IEEE, April 2001.

[7] Santosh P.Abraham and Anurag Kumar, "A New Approach for Asynchronous Distributed Rate Control of Elastic Sessions in Integrated Packet Networks," *IEEE/ACM Transactions on Networking*, February 2001.

[8] M. Vojnovic, J.-Y. Le Boudec, and C. Boutremans, "Global Fairness of Additive Increase and Multiplicative Decrease with Heterogenous Round Trip Times," in *Proc. IEEE Infocom 2000*, 2000.

[9] Steven Low, Larry Peterson, and Limin Wang, "Understanding TCP vegas: A Duality Model," in *Proc. SIGMETRICS 2001*, 2001, paper obtained from http://www.cs.princeton.edu/nsg/publications.html.

[10] F.P. Kelly, A.K. Maulloo, and D.K.H. Tan, "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Operations Research*, vol. 49, 1998.

[11] Jeonghoon Mo and Jean Walrand, "Fair End-to-End Window-Based Congestion Control," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556–567, October 2000.

[12] Steven H. Low and David E. Lapsley, "Optimization Flow Control-I: Basic Algorithm and Convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, December 1999.

[13] Aditya Karnik and Anurag Kumar, "Performance of TCP Congestion Control with Rate Feedback: Rate Adaptive TCP (RATCP)," in *Proc. IEEE Globecom 2000*, 2000.

[14] J. W. Roberts, "Engineering for quality of service," *K. Park and W. Willinger, editors, Self-Similar Network Trac and Performance Evaluation. Wiley-Interscience*, 2000.

[15] D. P. Heyman, T. V. Lakshman, and A. L. Neidhardt, "A new method for analyzing feedback-based protocols with applications to engineering web traffic over the internet," *Proc. ACM-SIGMETRICS'97*, 1997.

[16] Arthur W. Berger and Yaakov Kogan, "Dimensioning Bandwidth for Elastic Traffic in High Speed Data Networks," *IEEE/ACM Trans. Networking*, 2000.

[17] A.A.Kherani and Anurag Kumar, "Performance analysis of TCP with nonpersistent sessions," in *Workshop on Modelling of Flow and Congestion Control, INRIA, Ecole Normale Superieure, Paris*, September 2000.

[18] Claudio Casetti and Michela Meo, "A New Approach to Model Stationary Behavior of TCP Connections," *Proc. IEEE INFOCOM'2000*, April 2000.

[19] James Roberts and Laurent Massoulie, "Bandwidth sharing and admission control for elastic traffic," October, 1998.

[20] A. A. Kherani, "Rate Based Congestion Control for Ephemeral Best-Effort Sessions in Packet Networks," Master of Engineering Thesis, Indian Institute of Science, Bangalore, Jan 1999.

[21] Vern Paxson and Sally Floyd, "Wide-Area Traffic: the Failure of Poisson Modeling," *IEEE/ACM Trans. on Networking*, vol. 3, no. 3, pp. 226–244, 1995.

[22] E. Gelenbe and L. Mitrani, *Analysis and Synthesis of Computer Systems*, Academic Press, 1980.

[23] Pinaki Shankar Chanda, "Models for Traffic Engineering of Packet Networks with Non-Persistent Elastic Flows," Master of Engineering Thesis, Indian Institute of Science, Bangalore, Jan 2001.

[24] Dimitri Bertsekas and Robert Gallager., *Data Networks*, Prentice Hall, 1992.

[25] Anja Feldmann and Ward Whitt, "Fitting mixtures of exponentials to long-tail distributions to analyze network performance models," Tech. Rep., AT&T Laboratory - Research, December 16, 1996.

[26] Pinaki S. Chanda, Anurag Kumar, and Arzaad A. Kherani, "An Approximate Calculation of Max-Min Fair Throughputs for Non-Persistent Elastic Sessions," Extended Draft, can be downloaded from http://ece.iisc.ernet.in/~ anurag/.