FIT – Future Internet: Traffic Handling and Performance Analysis

Jorma Virtamo
Networking Laboratory, HUT
P.O.Box 3000, 02015 HUT, Finland
jorma.virtamo@hut.fi

Ilkka Norros
VTT Information Technology
P.O. Box 1202, 02044 VTT, Finland
ilkka.norros@vtt.fi

ABSTRACT

The FIT project has addressed a number of problems arising in controlling the traffic and providing quality of service in the Internet and in analyzing the performance of the system. An efficient recursive algorithm has been developed for calculating the flow level performance of elastic traffic under a bandwidth sharing scheme called balanced fairness (BF). The notion of BF has been extended and applied to the flow throughput estimation in ad hoc networks; also other applications of BF are being explored. Scheduling mechanisms, in general, and adaptive scheduling mechanism with delay bounds, in particular, were studied. A simple adaptive and distributed load balancing mechanism has been suggested and analyzed. In this system, traffic is gradually redistributed based on measured link loads, leading to a nearly optimal performance. Analytical results have been obtained on the performance of MAC protocols in ring networks employing optical burst switching. For traffic matrix estimation, the Gravey-Vaton method has been analyzed in detail in the case of Gaussian traffic variations. The stability problem of an overloaded network with measurement-based admission control has been analytically solved.

I. INTRODUCTION

FIT is a joint project between the teletraffic theory group of the Networking Laboratory of Helsinki University of Technology (HUT) and the VTT/UH (University of Helsinki) group. The site http://www.netlab.hut.fi/tutkimus/fit/ provides general information about the project. The project started in 2001 and continues to the end of 2004. For various reasons, the center of the gravity of the project at HUT has shifted towards the end of the project and at the time of this writing (March 2004) the project is still going on with five researchers working full time. Therefore, this report is not complete but more results will be produced.

At a general level, the objective of the project has been to study methods for traffic handling in the Internet in order to avoid or manage congestion and methods for providing Quality of Service (QoS) and QoS differentiation and, finally, to develop methods for analyzing the performance of such systems. The specific topics studied are detailed below.

Balanced Fairness (BF) discussed in Section II can, on one hand, be viewed as an ideal bandwidth sharing scheme. On the other hand, BF can be seen as a computational approximation tool for analyzing more easily implementable schemes like max-min fairness. Scheduling, which is the topic of Section III, is the key mechanism for QoS differentiation in the routers; packets belonging to different traffic classes are handled differently. Load balancing is one of the basic tasks of Traffic Engineering (TE) in a network. Section IV introduces a simple adaptive mechanism by which an almost ideal load distribution is obtained without knowing the traffic matrix. At the transport level, modern broadband networks utilize optical technology, which poses different type of traffic handling problems. In Section V, the performance of MAC protocols in an optical ring network using optical burst switching is analyzed. Many traffic handling operations at the network level require knowledge of current traffic demands between different origin-destination pairs. Estimating the traffic matrix on the basis of link load measurements, however, is a strongly under determined problem and poses a big challenge. Results of a study related to a recently introduced method are discussed in Section VI. Some traffic measurements have also been done in the project as briefly documented in the next section. Section VII deals with another traffic handling method, measurement based admission control for networks subject to overload. The analysis of a simple network model demonstrates how the lack of complete state information makes the stability of the system very sensitive to its parameters

II. BALANCED FAIRNESS

In the Internet, the bandwidth resources of the network are shared between all concurrent traffic flows. Typically, the sharing is determined by a flow control protocol such as the TCP (Transmission Control Protocol), which is used by the majority of applications. An objective of bandwidth sharing is some kind of fairness, but what fairness means when there are many resources (link capacities) shared by several flows with different routes is not obvious. Many definitions have been developed, including the classical max-min fairness and proportional fairness. These are special cases of a more general class of so-called utility based fairness criteria [1].

A new notion of fairness, called balanced fairness (BF), has recently been introduced by Bonald and Proutière [2]-[4]. The most important property of BF is that it leads to completely insensitive network performance: the performance under BF depends solely on the traffic loads on different routes but not of any other traffic characteristics (e.g., flow size distribution). This is a very desirable property from the point of view of networks design and management. Another remarkable property of BF is that the performance in a real dynamic setting, where new flows arrive stochastically and depart upon completion, can be evaluated in an analytic form for many important network topologies.

The analysis of a network performance under Balanced Fairness is based on the theory of Whittle queueing networks [5], i.e. networks of Processor Sharing (PS) nodes the capacities of which, $\phi_i(x)$, depend on the global network state x in a "balanced way". This condition means that there exists a balance function $\Phi(x)$ such that

$$\phi_i(x) = \frac{\Phi(x - e_i)}{\Phi(x)} \qquad \forall i, x, \tag{1}$$

where $x - e_i$ denotes the network state with one class-i customer less than in state x. Balanced Fairness refers to a capacity allocation that satisfies the balance condition (1) and at the same time utilizes the network resources maximally.

An important contribution of the FIT project was the development of a recursive method for an exact calculation of the normalization constant of a BF system [6, 7] (from the normalization constant all the performance metrics can be easily derived). Notably, the algorithm completely avoids solving the balance function explicitly. The recursion is also very efficient: in a system with n flow classes only 2^n numbers have to be recursively calculated, in contrast to some N^n numbers in a direct calculation, where N typically, depending on the desired accuracy, is of the order of 100.

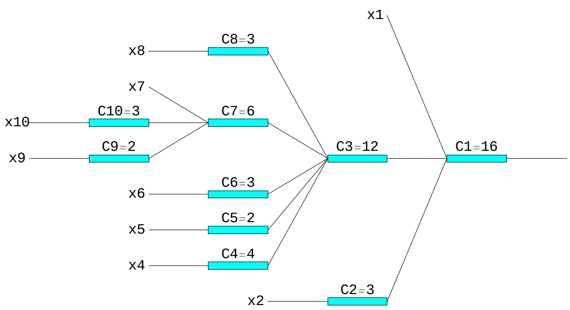

Figure 1: An example of a 4-level tree.

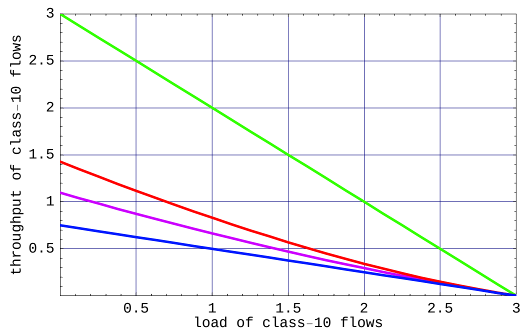

Figure 2: Comparison of class-10 throughput lower and upper bounds with the exact result. From bottom up: store-and-forward, parking lot, exact, deterministic.

The method has been applied to practically important topologies for access networks, like parking lot and general tree with and without access rate limits. An implementation of the recursion for a general tree network has been added to the Qlib library [8]. As an important side result it was proven in [7] that BF is Pareto efficient in all tree networks, i.e. no network resources are wasted by the balance condition. Explicit results have been worked out for several example networks. In Figure 1 a four level tree network is illustrated and Figure 2 shows the flow throughput in this network for class-10 flows going through the links 10, 7, 3 and 1.

As noted before, the distinguishing feature of BF is its insensitivity. In contrast, all utility based fairness schemes are sensitive except for some special network topologies, where they coincide with BF. Indeed, Bonald and Proutière have shown that in order for any queueing network of PS nodes to be insensitive the resource allocation has to be balanced. This result was extended to any network of symmetric queues in [9].

An extensive study of the sensitivity properties of various schemes was undertaken in [10]. The study confirms that the non-BF schemes are sensitive (their performance was evaluated by simulations). The sensitivity on the flow size distribution, however, is not very strong; more pronounced is the sensitivity with respect to so-called time scale changes. It was also found that, largely speaking, BF provides a reasonably good approximation for the performance of max-min fairness. Furthermore, it was confirmed that in a network topology called hypercycle, BF is not Pareto efficient but some capacity is wasted. This is illustrated in Figure 3. Fortunately it turns out that in a dynamic setting the impact of the capacity waste is minor. An analytical representation for the balance function of a symmetric 3-link hypercycle was obtained in [11].

The concept of BF was extended in [12] to the case where the constraints are more complicated than fixed link capacity constraints. In particular, [12] outlines an approach for the case where the flows can be split over several routes

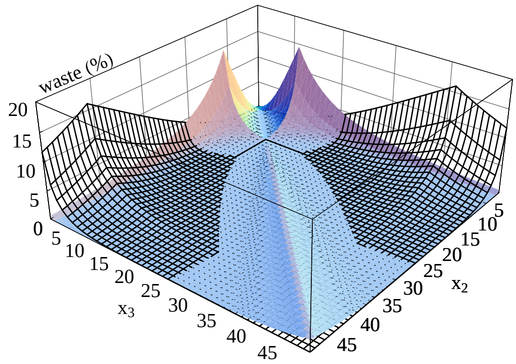

Figure 3: The capacity waste in a 3-link hypercycle.

and the constraints are then of the type encountered in multicommodity flow problems. The important insight is that in a given state x, given the balance function for all states $y < x$, the balance condition (1) fixes the capacity allocations up to a multiplicative factor $1/\Phi(x)$. Then, $\Phi(x)$ can be recursively determined by making it at each point as small as allowed by whatever constraints the system is subject to.

This idea was further developed in [13], where the joint problem of scheduling and resource sharing was studied in an ad hoc network, where the simultaneous use of some links is limited by the interference. In an idealized model, the effect of scheduling can be viewed as allowing, under certain limits, a shift of capacity from one link to another one. Again, the system constraints are not just fixed link constraints. By applying the extended BF principle, it was possible to determine both the schedule and resource sharing in such a way that the resources are maximally utilized while preserving the balance property guaranteeing the insensitivity and robustness of the system. In addition, analytical tractability is retained, at least for smaller systems.

In the FIT project, several studies related BF are still going on. In [14] the possibility of insensitive adaptive routing is being studied. In insensitive routing, not only the capacity allocations are balanced but also, separately or jointly, the arrival rates determined by the probabilistic routing. In another study [15], queueing network models, such as BF, are being applied to the analysis of flow throughputs of the data traffic in a cell of a cellular network with finite user population. Application of BF to closed queueing networks was explored in [16]. A general tutorial on Balanced Fairness is also available on the web site of the project [17].

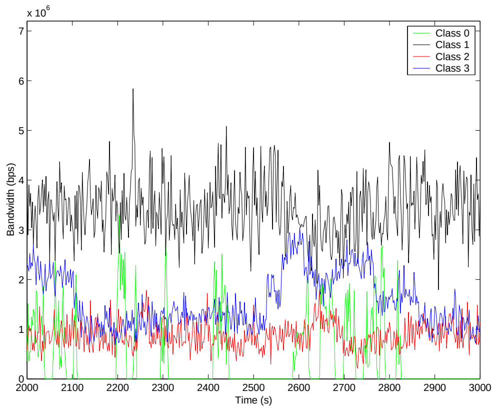

Figure 4: Bandwidth allocations as a function of time in a simulation run with an adaptive scheduler.

III. Scheduling and Quality Differentiation in DiffServ

During the last decade the Internet has developed into a public multiservice network that should be able to support heterogeneous applications and customers with diverse requirements. For this reason, quality of service (QoS) provisioning in the Internet has gained increasing attention. General service architectures, e.g. Differentiated Services (DiffServ) [18], have been proposed in the literature for providing QoS. In DiffServ, common resources are allocated among service classes. Packet scheduling is the mechanism primarily responsible for the allocation.

In [19] and [20], we studied issues related to quality differentiation, traffic mapping and scheduling in DiffServ. The focus was on the differentiation of two important quality parameters, capacity and delay. Two models were concerned in detail, absolute capacity differentiation and proportional delay differentiation with delay bound, and various packet schedulers were investigated by simulations. According to the results, provisioning and differentiation with static resource allocation methods is problematic but with some adaptive schedulers tunable so that consistent differentiation can be achieved. Based on these observations, we suggested that in DiffServ networks an adaptive scheduler with delay bound should be used for resource allocation. We also argued that from applications point of view it is beneficial to map different traffic types into separate service classes.

Figures 4 and 5 illustrate the behaviour of an adaptive scheduler (HPD with delay bound) in a simulation run with mixed traffic in four service classes. It can be seen that the bandwidth allocation follows quite well the development of queue lengths: if the queue of a class starts to build up, more resources are allocated to the class.

IV. Adaptive Load Balancing Using MPLS

Multiprotocol Label Switching (MPLS) [21] brings up new possibilities to improve the performance of IP networks. Notably the explicit routing of MPLS facilitates balancing the load by moving traffic from a congested part of the network to some other part in a well controlled way.

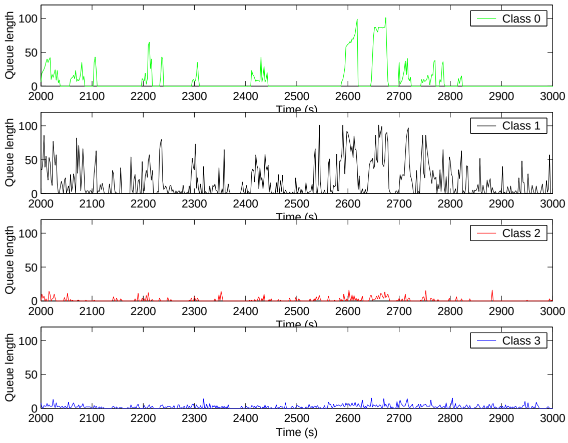

Figure 5: Queue lengths as a function of time in a simulation run with an adaptive scheduler.

In [22], we studied adaptive load balancing based on measured link loads without knowledge of the full traffic matrix. As an objective for load balancing we used minimizing the maximum link utilization and minimizing the mean delay in the network. A simple adaptive and distributed algorithm was presented to obtain these objectives. In addition, a numerical method was developed to evaluate the performance of the algorithm. The method was applied to three test networks. The results showed that the maximum link utilization and the mean delay obtained in a reasonable number of iterations are very close to the optimal values, even with random traffic fluctuations in the time-scale of the measurements.

Figure 6 illustrates how our adaptive algorithm approaches the lowest possible level of the maximum link utilization (0.54) in less than 100 iterations when applied to a test network with 10 nodes, while the standard min-hop routing results in a much higher level of congestion corresponding to the maximum link utilization of 0.88. Three different adaptation granularities are applied. With a finer granularity (g = 50, 100), the adaptive algorithm converges slightly slower but much more smoothly than with a coarse granularity (g = 10).

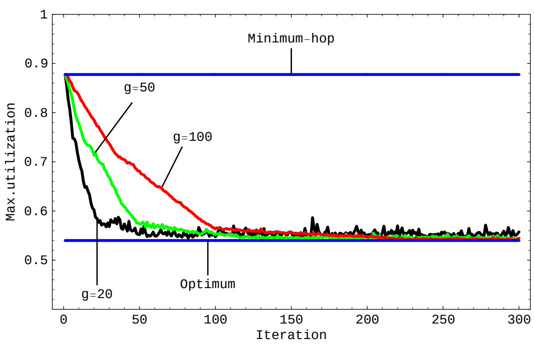

Figure 6: The maximum link utilization as a function of the number of iterations in a numerical evaluation run of an adaptive load balancing algorithm.

V. Optical Burst Switching

Wavelength division multiplexing (WDM) is a technique, where several optical channels are used concurrently to transfer massive amounts of information, even terabits per second, in a single optical fibre. Optical networks employing WDM play already an important role in the current backbone networks and the trend is that WDM will be adopted also in MAN and finally in LAN networks.

In optical burst switching network optical bursts are used to transfer the data [25]. Each burst consists of several concatenated (IP) packets all having the same destination node and thus routed along the same path. Hence, the optical burst switching (OBS) can be seen as an intermediate step from the wavelength routed networks (i.e. circuit switching) towards the optical packet switching.

In an OBS network the necessary resources are typically reserved only for the duration of the burst. In particular, the source node first sends a control packet or frame to inform the receiver (and possibly intermediate nodes) about the coming burst. The receiver and intermediate nodes each check, one at a time, if the switches along the path can be configured to deliver the burst successfully towards its destination. Meanwhile, after a certain offset time, the actual data burst is sent along the same path without waiting for any acknowledgment from the receiver (or intermediate nodes).

Optical ring network is a suitable solution for metropolitan area networks (MAN). One cost effective solution using the OBS paradigm is described in [26, 23], where each station of the ring has a dedicated fixed “home wavelength channel” for transmitting its bursts. In addition to these data channels also a shared control channel with a certain number of circulating control frames is used to inform the other nodes about the arriving bursts resulting a time slotted system. Thus, the number of wavelength channels W is equal to $N+1$, where N is the number of stations. Furthermore, each station has only one adjustable receiver. Thus, no transmissions collide in the fibre. But, as each station can listen to at most one channel at a time, burst losses may occur at the receiver in the case two or more concurrent bursts have the same destination node. The traffic pattern for a case with $N = 4$ nodes is illustrated in Figure 7.

In [24] we have studied, both analytically and numerically, the performance of the MAC protocols proposed in [26] under static traffic conditions. In particular, we have considered both random order and round-robin transmission policies. When a station operates in random order policy the transmission queue to be served next is chosen randomly

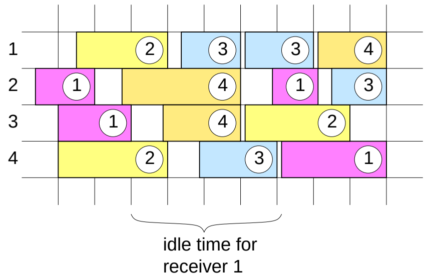

Figure 7: Arrival process in OBS-ring network.

among the non-empty queues. In round-robin transmission policy the non-empty transmission queues are served in a fixed order so that during a one full period one burst is transmitted from each queue. In the analysis we have considered an arbitrary receiver and derived formulae for the burst blocking probability and the so called (receiver) efficiency, i.e. the proportion of the time the receiver is active. In Figure 8 the efficiency is depicted for both random order and round-robin transmission policies in a certain traffic case. It can be seen that the round-robin order, while generally being a more fair, leads to a worse efficiency, and hence, to a higher blocking probability.

In [24], special attention has been given to the performance under an extremely heavy load, where each source has always bursts to be sent to all $N-1$ other destinations, i.e. the offered load is $\rho = 1$. Although this symmetric heavy traffic scenario does not hopefully exist, it gives a lower bound on blocking probability for each MAC protocol and allows comparing their worst case performances. Note that in an ideal situation the blocking probability would be zero and each receiver busy all the time, leading to an average pairwise throughput of $C/(N-1)$, where C is the capacity of one channel and N the number of nodes. However, in [24] we have shown that without any coordination between the nodes the actual throughput will be considerably less, i.e. about a half of that. This heavy traffic scenario

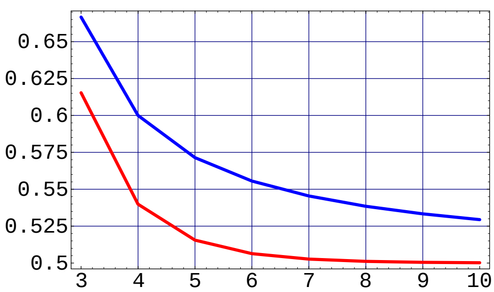

Figure 8: Receiver efficiency under a heavy traffic load for random order (upper curve) and round-robin (lower curve) transmission policies as a function of the number of active nodes.

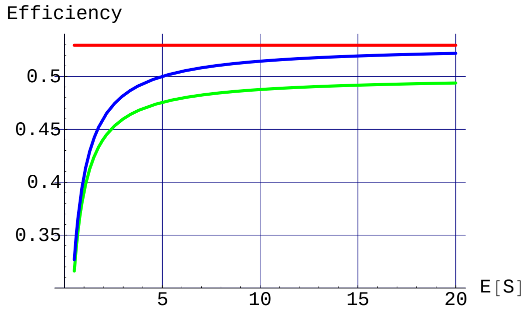

Figure 9: Efficiency for $N = 10$ (upper curve) and $N = \infty$ (lower curve) as a function of mean burst size. The constant line represents the efficiency for very long bursts.

is depicted in Figure 9. From the figure it can be seen that the mean burst size should be about 5 slot times or more in order to achieve a reasonably high efficiency, where the slot time is equal to the interarrival time of control frames.

VI. Traffic Matrix Estimation

The traffic matrix $\boldsymbol{x}$, which gives the volume of traffic between each origin/destination (OD) pair in the network, is a required input in many network management tasks. Unfortunately, the traffic matrix cannot be directly measured. Only the link loads $\boldsymbol{y}$ and the routing matrix $\boldsymbol{A}$ are available. These satisfy the relation

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}. \quad (2)$$

Since in any realistic network there are many more OD pairs than links, the problem of solving $\boldsymbol{x}$ from $\boldsymbol{A}$ and $\boldsymbol{y}$ is highly underdetermined and the problem is ill-posed.

To overcome the ill-posedness, some type of additional information has to be brought in to solve the problem. Typically, prior information is incorporated into the traffic matrix estimation using different traffic models. Most promising proposed methods include Bayesian inference techniques and network tomography [28, 29, 30]

The Vaton-Gravey iterative Bayesian method [31] consists of iteration and exchange of information between two "boxes"; the first calculating an estimate for the traffic matrix, given the observed link counts and parameter values of the prior, and the second updating the parameter values. Both boxes involve extensive numerical simulations or numerical algorithms.

In the FIT project, a study [27] of the above idea has been conducted with simplifying assumptions. The aim has been to gain insight into the method and, in particular, the output of the first box by examining a model simple enough to be computed analytically. Independence and normality are assumed for the OD pair distributions. This reduces the complexity of the Vaton-Gravey method. Our prior information consists of the mean and the covariance matrix of the Gaussian distribution. The attractive feature of this approach is that the distribution conditioned on the link counts is again Gaussian and analytical results can be obtained. When iteration is performed, it turns out that the expected value of the mean does not change in the iteration after the first conditioning on link counts has been made.

We have shown that when the mean, $\boldsymbol{m}$, and the covariance, $\boldsymbol{C}$, of the Gaussian distribution is iterated, the expectation of the result $(\boldsymbol{m}, \boldsymbol{C})$ of an iteration round can be written as

$$\begin{cases} \boldsymbol{m}^{(i+1)} &= \boldsymbol{GA\mu} + \boldsymbol{H}\boldsymbol{m}^{(i)}, \\ \boldsymbol{C}^{(i+1)} &= \tilde{\boldsymbol{C}} + \boldsymbol{GA\Sigma}(\boldsymbol{GA})^{\mathrm{T}}, \end{cases}$$

where $\boldsymbol{G}$, $\boldsymbol{H}$ and $\tilde{\boldsymbol{C}}$ depend on the covariance of the prior distribution, while $\boldsymbol{A}$ and $\boldsymbol{\Sigma}$ are constant matrices. The index i refers to the iteration round. We have proven that the mean $\boldsymbol{m}$ does not change after the first iteration.

Figure 10 illustrates the situation when two OD pairs are to be estimated from a single link count. On the left is the prior distribution. The real distribution that we are trying to find out is shown on the right. From equation (2) we know the sum of the two variables x_1, x_2, telling that the mean of the distribution is somewhere along the line shown in the picture. The resulting distribution is shown in red, and is the estimate of the real distribution based on the prior distribution and link loads. The green dots are produced by sampling from the conditional distribution, and are seen to coincide with the explicitly calculated result, as expected.

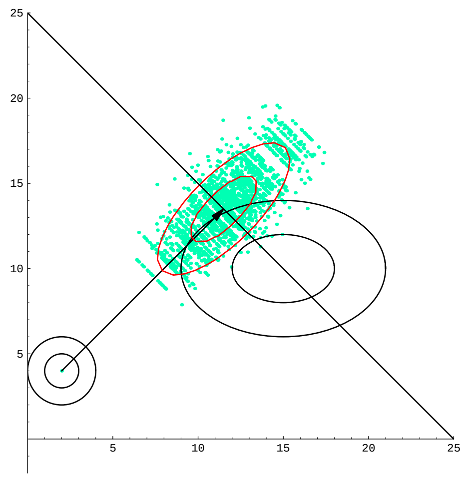

Figure 10: Result of the first iteration.

Directions for future work may include for example studying the possibility of utilizing the covariance matric in the estimation. This requires assuming a certain relation between the mean and the variance of the OD pairs.

VII. Modem Pool Traffic Survey

We continued collection and analysis of the HUT modem pool statistics started in the previous project (Com^2) funded by the Academy of Finland. In addition to the earlier samples from years 1997 and 1998 surveyed in [35], we got new samples in 2001 and 2002. A comparative study was carried out and the results were presented in [36]. An illustration of the results for 2002 is given in Figure 11. A comparison with the results of 2001 shows that the daily usage profile has remained very much the same but the overall traffic level has decreased almost to a half, presumably due to a shift from modem to ADSL connections.

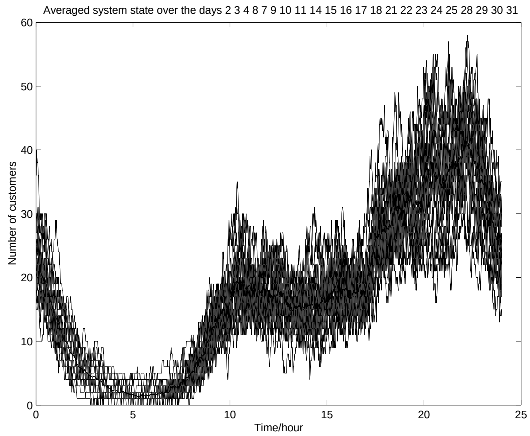

Figure 11: The number of student users as a function of time during the weekdays in October of 2002.

VIII. Stability Analysis of an Admission Control Mechanism

In DiffServ Networks, the suggested resource allocation mechanisms for bandwidth brokers are based on incomplete information about the network state. In this project, we have investigated the impact of this lack of information to the stability properties of the mechanisms. To facilitate analytical treatment, we have restricted ourselves to a simple two-server network with feedback admission control, depicted in Figure 12.

Rather surprisingly, it was found that this simple network has a non-trivial stability region. A paper on this discovery has been submitted for publication [32]. The solution of the problem is depicted in Figure 13, with unit input rate. The parameter space of the server rates is partitioned into four regions. Here A_1 is the domain where the system is never stable, no matter how strict admission control is employed. Regions A_2 and A_3 represent the situation where the overloaded network can be stabilized; A_2 requires strict enough admission control to be stable, while A_3 is stable with any

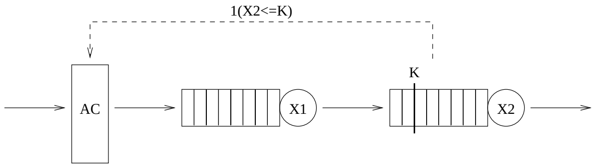

Figure 12: The admission controller (AC) tracks the state of the server 2. If the size of the queue for the server 2 exceeds a given threshold K, the AC will block all incoming traffic.

control threshold. In A_4 there is no overload present, and thus the network is always stable. The phase diagram illustrates how accelerating server 1 drives the system towards more stable regions, while rather paradoxically, speeding up server 2 may unstabilize the network. In [32] these results are proven and extended to more complex controllers, giving easily verifiable stability conditions for the network in terms of the system parameters.

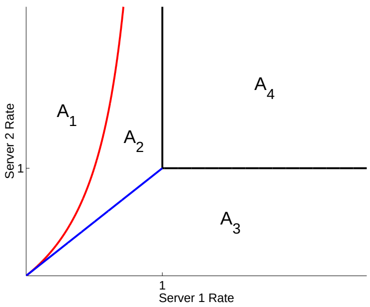

Figure 13: Phase diagram of the network with feedback admission control.

Another related problem studied in this project was connected with traffic modelling. An interesting new limit process had been recently found (later published as [33]) that resulted in simultaneous increase of the number of sources and the timescale. A proof of an analogous result for univariate distributions with a somewhat different traffic model was found independently, but the authors of [33] extended their paper to cover this model as well. The attention was moved to another aspect of long-range dependent traffic: its behaviour under measurement-based admission control. This means a scheme where traffic flows are accepted to the network according to the measured load. The goal was, in rigorous mathematical terms, to understand the empirical observation of [34] that such a control almost removes long-range dependence, and to study how reliable this kind of feedback mechanisms would be. This type of question is closely connected with the question studied in [32], and there are plans to extend the type of results obtained in [32] for stability of Markovian networks to limiting distributions with heavy-tailed input processes.

IX. Results and Impacts

Within the FIT project so far two Master's theses have been completed (J. Antila and V. Timonen) and one more will be completed in 2004 (S. Liu). One PhD work will be completed in 2004 (E. Hyytiä), two others are at an intermediate phase (L. Leskelä and R. Susitaival), and two more at an early phase (I. Juva and J. Leino).

The FIT project has been an important element in helping the HUT and VTT groups to strengthen their international contacts and collaboration. During the project the groups have been actively participating in the action COST279, where also work done in this project has been reported.

Both the VTT and HUT groups were invited to become members in Euro-NGI, a Network of Excellence of the EU 6th Framework Programme. Altogether 58 partners are involved in this three-year activity, which was started in December 2003.

The appropriation for Senior Scientist received by Prof. Virtamo in association to this project enabled very fruitful visits to France Telecom R&D and to Cambridge University. The work on balanced fairness was initiated during these visits, and the collaboration with the researcher's at France Telecom R&D continues. Several studies related to this topic are going in the Networking Laboratory.

Prof. Virtamo has served as a member of Scientific Committee of the forthcoming program on Queueing Theory and Teletraffic Theory at the Institute Mittag-Leffler of the Royal Swedish Academy of Sciences. Several researcher of the project are going to participate in the Program by a longer stay at the Institute during autumn 2004.

Networking Laboratory of HUT hosted the 16th Nordic Teletraffic Seminar, NTS-16, which was held in Otaniemi in August 2002 [37]. The seminar was supported by a separate grant from the Academy of Finland.

The results of the FIT project have been annually presented in a half-day seminar jointly organized with the domestic COST279 project funded by Tekes. Researchers from the academia and industry have been invited to attend the seminar. Links to the programs of the last two seminars are given in [38].

The program library of teletraffic theory functions [8] has been maintained and developed. New functions and algorithms have been added into the library, notably those related to calculating the performance under balanced fairness, as well as algorithms needed for network calculations and synthetic network generation for simulation purposes.

References

[1] J. Mo and J. Walrand, 2000, "Fair end-to-end window-based congestion control", IEEE/ACM Transactions on Networking, Vol. 8, pp. 556-567.

[2] T. Bonald, A. Proutière, 2003, "Insensitive bandwidth sharing in data networks", Queueing Systems, Vol. 44, pp. 69-100.

[3] T. Bonald, A. Proutière, 2002, "Insensitivity in processor-sharing networks", Performance Evaluation, Vol. 49, pp. 193-209.

[4] T. Bonald, A. Proutière, 2004, "On performance bounds for balanced fairness, Performance Evaluation", Vol. 55, pp. 25-50.

[5] R. Serfozo, 1999, "Introduction to Stochastic Networks", Springer-Verlag, Berlin.

[6] T. Bonald, A. Proutière, J. Roberts and J. Virtamo, 2003, "Computational aspects of balanced fairness", Proceedings of the 18th International Teletraffic Congress (ITC-18), pp. 801 - 810.

[7] T. Bonald and J. Virtamo, 2004, "Calculating the flow level performance of balanced fairness in tree networks", to appear in Performance Evaluation.

[8] Qlib, Program Library for Traffic Theory Functions, http://www.netlab.hut.fi/qlib/.

[9] J. Virtamo, 2004, "Insensitivity of a network of symmetric queues with balanced service rates", submitted.

[10] V. Timonen, 2003, "Simulation Studies on Performance of Balanced Fairness", Master's Thesis, Helsinki University of Technology, Networking Laboratory, Report 6/2003.

[11] J. Virtamo, 2003, "On the balance function of a homogeneous hypercycle", Technical Report.

[12] L. Massoulié and J. Virtamo, 2003, "Extending the notion of balanced fairness", Technical Report.

[13] A. Penttinen and J. Virtamo, 2004, "Performance of wireless ad hoc networks under balanced fairness", to be presented in Networking'04, Athens.

[14] J. Leino, 2004, "Insensitive Routing", Tech. Report.

[15] S. Liu, 2004, "Quasi stationary limit of the data traffic performance of a wireless cell with finite user population", work in progress (Master's thesis).

[16] J. Virtamo, 2003, "Calculating the normalization constant for closed networks", Technical Report.

[17] J. Virtamo, 2003, "Balanced fairness – insensitive bandwidth allocation", tutorial presentation, 80 pages.

[18] S. Blake, D. Black,M. Carlson, E. Davies, Z. Wang and W. Weiss, 1998, "An Architecture for Differentiated Services", IETF RFC 2475.

[19] J. Antila, 2003, "Scheduling and quality differentiation in Differentiated Services", Master's Thesis, Helsinki University of Technology, Networking Laboratory, Report 5/2003.

[20] J. Antila and M. Luoma, 2003, "Scheduling and quality differentiation in Differentiated Services", Proceedings of MIPS, November 2003, pp. 119-130.

[21] E. Rosen, A. Viswanathan and R. Callon, 2001, "Multiprotocol Label Switching Architecture", IETF RFC 3031.

[22] R. Susitaival, S. Aalto and J. Virtamo, 2004, "Adaptive load balancing using MPLS", submitted for publication.

[23] T. Battestilli, 2002, "Optical burst switching: A survey", Technical Report TR-2002-10, North Carolina State University.

[24] E. Hyytiä and L. Nieminen, 2004, "Analytical loss models for MAC protocols in optical ring network operating under a static traffic load", submitted for publication.

[25] C. Qiao and M. Yoo, 1999, "Optical burst switching (OBS) - a new paradigm for an optical internet", Journal of High Speed Networks, Vol. 8, pp. 69-84.

[26] L. Xu, H.G. Perros, and G.N. Rouskas, 2002, "A simulation study of protocols for optical burst-switched rings", Proceedings of Networking 2002, pp. 863-874.

[27] I. Juva, P. Kuusela, J. Virtamo, 2004, "A Case Study on Traffic Matrix Estimation under Gaussian Distribution", Technical Report.

[28] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, 2002, "Traffic matrix estimation: Existing techniques and new directions", ACM SIGCOMM 2002, Pittsburg, USA.

[29] C. Tebaldi and M. West, 1998, "Bayesian inference on network traffic using link count data", Journal of the American Statistical Association, Vol. 93, pp. 557-576.

[30] Y. Vardi, 1996, "Network Tomography: Estimating Source-Destination Traffic Intensities from Link Data", Journal of the American Statistical Association, Vol. 91, pp. 365-377.

[31] S. Vaton, A. Gravey, 2004, "Network tomography: an iterative Bayesian analysis", Proceedings of the 18th International Teletraffic Congress, ITC-18.

[32] L. Leskelä, 2004, "Stabilization of an Overloaded Queueing Network Using Measurement-Based Admission Control", submitted for publication.

[33] R. Gaigalas and I. Kaj, 2003, "Convergence of scaled renewal processes and a packet arrival model", Bernoulli, Vol. 9, No. 4, pp. 671-703.

[34] G. Bianchi, V. Mancuso and G. Neglia, 2002, "Is Admission-Controlled Traffic Self-Similar?", Proceedings of Networking 2002, Pisa, Italy.

[35] J. Lakkakorpi, 1999, "Traffic in modem pools of Helsinki University of Technology", Helsinki University of Technolgy, Laboratory of Telecommunications Technology, Technical Report.

[36] A. Matinlauri, 2003, "Modem pool traffic survey continued", Helsinki University of Technolgy, Networking Laboratory, Technical Report.

[37] http://www.netlab.hut.fi/nts16/.

[38] http://www.vtt.fi/tte/tte21/cost279/seminar.htm, http://www.netlab.hut.fi/tutkimus/cost279/FITseminar_2004.htm