A Distributed Scheme for Robust On-Line Network Engineering

Johannes Göbel
Department of Informatics
University of Hamburg
22527 Hamburg, Germany
aek1@cs.sun.ac.za

Anthony E. Krzesinski and Dieter Stapelberg
Department of Mathematical Sciences
University of Stellenbosch
7600 Stellenbosch, South Africa
aek1@cs.sun.ac.za

***Abstract*— This paper presents a model of an on-line network bandwidth management scheme: at random time instants bandwidth prices are computed and are used to adjust the bandwidths of the network paths in response to the traffic conditions. The distinguishing features of the scheme are that it works without centralised control and thus scales to large networks, and that rather than using traffic engineering to move network flows to where the network bandwidth is located, network engineering is used to move bandwidth to where the flows are located. Simulation results show that the reallocation scheme provides prompt and robust bandwidth provisioning.**

Keywords: bandwidth prices; bandwidth reconfiguration; distributed control; network planning and optimisation; scalability.

I. INTRODUCTION

This paper presents a model of network bandwidth management that is based on a distributed scheme for automatic bandwidth reallocation [3]. The reallocation scheme can be used in the context of a connection-oriented network where relatively long-lived paths are used to provision resources for connections (the terms path/route and connection/call will be used interchangeably), whose average holding times are much less than the path lifetimes. Multi-Protocol Label Switching (MPLS) networks [9], in which Label Switched Paths (LSPs) act as the long-lived paths, provide a possible environment in which such a bandwidth reallocation scheme could be useful.

The reallocation scheme is restricted to connection-oriented traffic with admission controls. Hence we use the terms connections and lost connections or equivalently calls and lost calls. The number of connections that can be simultaneously carried on a route depends on the amount of bandwidth allocated to the route. However, at any point in time, it is possible that due to traffic fluctuations or network equipment changes, the connections in service on one route are using only a small part of the bandwidth allocated to that route, while the bandwidth on another route is heavily utilised. Two methods have been developed to deal with this situation. First, the arriving calls can be offered to the least utilised route of a multi-path route connecting the origin-destination pair; second, bandwidth can be transferred from underutilised routes to overutilised routes.

The first method forms the basis of traffic engineering (TE) where network flows are moved to where the network bandwidth is located. Many TE methods have been developed to manage connectivity and to deliver resilience and Grade of Service (GoS) across the network. See [8] and the references therein for a description of TE methods and applications. The second method forms the basis of network engineering (NE) where bandwidth is moved to where the network flows are located. This paper investigates the use of NE to manage the network GoS.

A systematic way of applying NE is to assign a broker referred to as a bandwidth manager to each route. At random time instants the manager calculates the expected value, over a short period of time known as the planning horizon, of an extra unit of bandwidth (the buying price) and also the expected value, over the same short period of time, that the route would lose should it give up a unit of bandwidth (the selling price). Bandwidth can then be transferred from routes that place a low value on bandwidth to routes that place a high value on bandwidth. An essential component of the reallocation scheme is that the bandwidth transfer process is based on local information only.

The bandwidth managers are thus aware of local resource demands and bandwidth prices, and reallocate bandwidth among themselves in order to maintain the performance of their routes. In this way the managers are autonomous, act without centralised control from a system coordinator and behave entirely according to local rules. Such a scheme is distributed and scalable.

This paper evaluates the robustness of such a bandwidth reallocation scheme. The remainder of the paper is organised as follows. Section II presents a method to compute buying and selling prices for bandwidth in connection-oriented networks. Section III discusses how the prices are used in the bandwidth reallocation scheme. Section IV investigates the robustness of the reallocation scheme with respect to several potential uncertainties and/or inaccuracies in the parameter values which determine the operation of the reallocation scheme. Section V evaluates how the reallocation scheme can be used to protect the performance of one traffic class from the demands of another traffic class. The conclusions are presented in Section VI.

II. THE PRICE OF BANDWIDTH

Consider a connection-oriented network where requests for connection on route r arrive in a Poisson stream. If the current number of connections on route r is less than the capacity of the route (the terms bandwidth and capacity are used interchangeably), the request is accommodated, else the request is rejected. The connection holding times are exponentially distributed.

We use a bandwidth pricing method [4] which models a route as an Erlang loss system and computes the expected lost revenue due to connections being blocked, conditional on the system starting in a given state. The expected lost revenue is used to compute both a buying price and selling price for bandwidth, relying on knowledge of the state at time zero.

A manager is assigned to each route and we assume that a route's manager is making reallocation decisions for a planning horizon of τ time units. We can regard the value of U extra units of bandwidth, or the buying price, as the difference in the total expected lost revenue over time $[0, \tau]$ if the route were to increase its bandwidth by U units at time zero. Likewise, we can calculate the selling price of U units of bandwidth as the difference in the total expected lost revenue over time $[0, \tau]$ if the route were to decrease its bandwidth by U units.

For a route with bandwidth C, let $R_{c,C}(\tau)$ denote the expected revenue lost in the interval $[0, \tau]$, given that there are c connections at time 0. The buying and selling prices $B_{c,C}(\tau, U)$ and $S_{c,C}(\tau, U)$ of U units of bandwidth are then given by

$$B_{c,C}(\tau, U) = R_{c,C}(\tau) - R_{c,C+U}(\tau)$$

and

$$S_{c,C}(\tau, U) = \begin{cases} R_{c,C-U}(\tau) - R_{c,C}(\tau), & 0 < c \leq C - U \\ R_{C-U,C-U}(\tau) - R_{c,C}(\tau), & C - U < c \leq C. \end{cases}$$

$R_{c,C}(\tau)$ is given by [4] the inverse of the Laplace transform

$$\widetilde{R}_{c,C}(s) = \frac{\theta(\lambda/s) P_c(s/\lambda)}{(s + C\mu) P_C(s/\lambda) - C\mu P_{C-1}(s/\lambda)}$$

where λ and μ are the parameters of the exponential connection arrival and connection service processes respectively, θ is the expected revenue earned per connection,

$$P_c(s/\lambda) = (-\mu/\lambda)^c \ \Gamma_c^{(\lambda/\mu)}(-s/\mu)$$

and

$$\Gamma_c^{(\lambda/\mu)}(-s/\mu) = \sum_{k=0}^{c} \binom{c}{k} \binom{-s/\mu}{k} \left(\frac{-\lambda}{\mu}\right)^{c-k} k!$$

is a Charlier polynomial [6]. A computationally stable method for efficiently calculating the functions $\widetilde{R}_{c,C}(s)$ is presented in [5]. These are numerically inverted using the Euler method [1] to yield the $R_{c,C}(\tau)$.

Fig. 1 presents the buying price $B_{c,C}(\tau, U)$ of bandwidth on a route of capacity $C = 100$. The planning horizon $\tau = P/\eta$ where the signalling rate $\eta = V\lambda$ is the rate at which reallocations are attempted. The planning ratio P is set to 1. The signalling ratio V is set to 0.1. The bandwidth reallocation unit U is set to four. The choice of these values is motivated in [3]. Fig. 1 shows that the buying price increases as the connection arrival rate λ and the initial occupancy c increase. The selling price $S_{c,C}(\tau, U)$ exhibits similar behaviour, and the prices are strictly convex functions of C is the sense that, for all values of U, $B_{c,C}(\tau, U) < S_{c,C}(\tau, U)$.

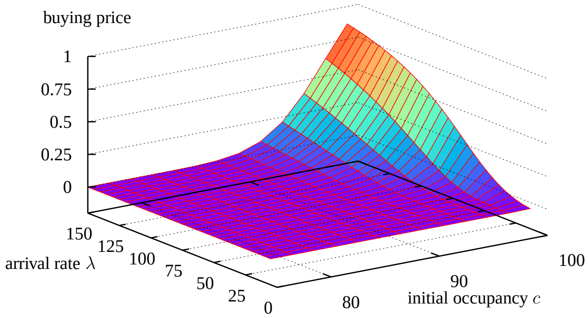

Fig. 1. The buying price of bandwidth as a function of the connection arrival rate λ and the initial occupancy c of the route: the planning horizon $\tau = 10/\lambda$.

III. A DISTRIBUTED BANDWIDTH REALLOCATION SCHEME

The buying and selling prices $B^{(r)}_{c_r,C_r}(\tau_r, U)$ and $S^{(r)}_{c_r,C_r}(\tau_r, U)$ of bandwidth on route r vary over time because the occupancy c_r varies over time. It is therefore likely that situations will arise where it will be advantageous to reallocate bandwidth among the routes.

As stated previously, the defining ingredient in our reallocation scheme is the manner in which bandwidth is reallocated among the routes. Bandwidth reallocation takes place between transit routes and their constituent direct routes. If only such "local" transfers are permitted, we avoid the need to compute and choose from the large set of possible reallocations and we do not need to consider the potentially widespread implications of a particular local reallocation.

The bandwidth reallocation scheme works as follows. At specified instants of time, a comparison is made between the buying price of bandwidth on transit route r and the sum of the selling prices on the constituent direct routes $\ell \in \mathcal{D}_r$, and between the selling price of bandwidth on transit route r and the sum of the buying prices on the constituent direct routes $\ell \in \mathcal{D}_r$. For a transit route r, if

$$B^{(r)}_{c_r,C_r}(\tau_r, U) > \sum_{\ell \in \mathcal{D}_r} S^{(\ell)}_{c_\ell,C_\ell}(\tau_r, U), \tag{1}$$

then the transit route acquires U units of bandwidth from each of its constituent direct routes. Alternatively, if

$$S^{(r)}_{c_r,C_r}(\tau_r, U) < \sum_{\ell \in \mathcal{D}_r} B^{(\ell)}_{c_\ell,C_\ell}(\tau_r, U), \tag{2}$$

then the transit route releases U units of bandwidth to each of its constituent direct routes. The above two inequalities cannot be simultaneously satisfied. If neither condition is satisfied, no reallocation occurs. More precisely, in order to acquire (release) bandwidth, the buying (selling) price of bandwidth on a transit route must be greater than (be less than) the accumulated selling (buying) price of bandwidth on the constituent direct routes by a small positive amount ϵ known as the *floor price*. This is done to prevent oscillating bandwidth reallocations: unless stated otherwise, $\epsilon = 10^{-6}$.

The manager of a transit route obtains the buying and selling prices by means of a signalling mechanism. Signals or control packets are sent at random intervals of time along each transit route, recording the buying and selling prices of the constituent direct routes. When a control packet reaches the egress router, a calculation is performed: if the sum of the direct route buying prices is greater than the transit route selling price, then the transit route will give up U units of bandwidth, which are taken up by each of the direct routes. Alternatively, if the sum of the direct route selling prices is less than the transit route buying price, then each of the direct routes will give up U units of bandwidth, which are taken up by the transit route. The control packet is returned along the reverse route from the egress router to the ingress router making the necessary reallocations to the capacities of the constituent direct routes. The capacity of the transit route is adjusted when the control packet reaches the ingress router.

See [3] for a discussion on how often the managers of transit routes send out control packets, the determination of the planning horizons τ_r used in the calculation of buying and selling prices and the determination of the amount U of bandwidth to be transferred in each reallocation. Signalling protocol issues such as the impact of signalling delays which will give rise to working with out-of-date bandwidth pricing information, and the impact of lost signalling packets which necessitates the refresh/cancellation of soft state are also discussed in [3].

IV. THE ROBUSTNESS OF THE REALLOCATION SCHEME

In this section we investigate the robustness of the reallocation scheme with respect to several potential uncertainties and/or inaccuracies in the parameter values which determine the operation of the reallocation scheme, namely the value of the floor price ϵ, the estimated values of the connection arrival rates λ_r, and the possibility that the bandwidth managers occasionally make incorrect reallocation decisions.

The bandwidth reallocation scheme is applied to the network model presented in Fig. 2 which represents a fictitious 29-node 45-link network model based on the geography of Europe. (The reallocation scheme has also been tested on networks models derived from Rocketfuel data [10].) Each origin-destination pair is connected by a single least cost (shortest hop count) route. (The reallocation scheme can readily be extended to the case where each (O,D) pair is connected by several routes.) The network model and its traffic matrix are described in [3]. A simulator was developed to model the connection arrival, admission and service processes, the signalling process, bandwidth price calculation and bandwidth reallocation. The simulation experiments assume that the connection arrival, the signalling and connection holding processes are exponential. Other holding time distributions were also investigated in [3].

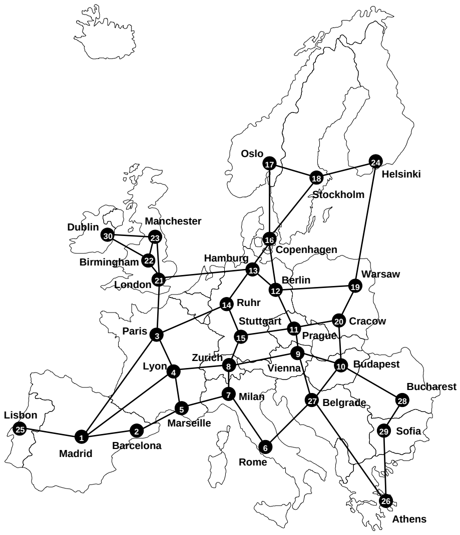

Fig. 2. The network model.

A. Floor price sensitivity

The floor price ϵ denotes the minimum reduction in the expected lost revenue that is necessary for a bandwidth reallocation to take place. Although an optimum value for the floor price can be determined experimentally, there would be no need to find a near-optimal value for the floor price if the reallocation scheme were relatively insensitive to the value of the floor price.

Fig. 3 presents the GoS, the revenue loss and the reallocation ratio for several values of the floor price ϵ. Confidence intervals are omitted since they are sufficiently small to be indistinguishable from the plot itself. The figure shows that the reallocation scheme is relatively insensitive to the value of the floor price. Note that a larger value for the floor price will reduce the number of successful reallocations. This will lead to a decrease in the reallocation success ratio which is defined as

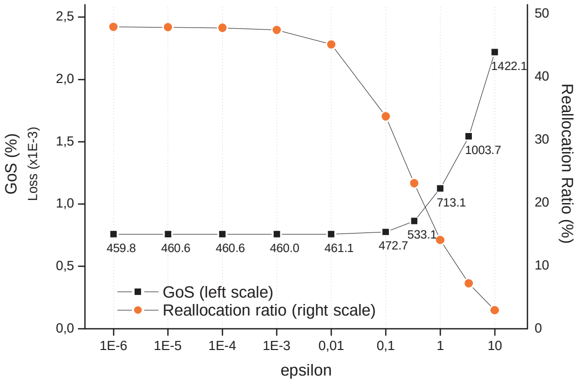

Fig. 3. The reallocation scheme is relatively insensitive to the floor price.

the number of successful reallocations versus the total number of reallocations attempted. This ratio quantifies the wasted effort spent in gathering pricing information for capacity reallocation attempts which, if they were implemented, would result in revenue gains which are less than the floor price.

The bandwidth reallocation scheme therefore does not depend on reallocations which yield marginal expected profits. This may be surprising since from an economic viewpoint a market equilibrium (a bandwidth assignment state that minimises loss thus clearing bandwidth supply and demand) potentially depends on conducting exchanges with marginal expected profit. Note however, that such a deferred trade may still be executed later when the demand of the party seeking to acquire bandwidth, and therefore the buying price, increases.

In addition, it is unlikely that the many accumulated small gains of the omitted bandwidth reallocations will have a significant effect on the GoS or on the revenue loss, because the number of the marginally profitably trades is small. With reference to Fig. 3, compare the slow decline in the reallocation success ratio for $\epsilon = 10^{-6}$ $(47.9\,\%)$ when increased up to 10^{-3} $(47.5\,\%)$ or even 10^{-2} $(45.2\,\%)$.

B. *Incorrect reallocation decisions*

Another aspect of the robustness of the reallocation scheme is the impact of incorrect reallocation decisions. For example, it cannot be guaranteed that all reallocation decisions conform to what is recommended by the reallocation scheme, since data transmission errors may result in corrupt, though plausible prices or may communicate incorrect reallocation instructions.

Incorrect reallocation decisions on transit route r are modelled by assuming that with probability P each reallocation attempt departs from its standard behaviour as follows

- when Eqn. (1) holds, instead of acquiring U units of bandwidth, bandwidth is released to the constituent direct routes if possible $(C_r - c_r > U)$, else no bandwidth reallocation is conducted
- when Eqn. (2) holds, instead of releasing U units of bandwidth, bandwidth is acquired from the constituent direct routes if possible $(C_{r'} - c_{r'} > U$ for each $r' \in \mathcal{D}_r)$ else no bandwidth reallocation is conducted
- when neither Eqn. (1) nor Eqn. (2) holds, instead of no bandwidth reallocation being executed, bandwidth is either acquired or released (if possible, see above) with equal probability, depending on the outcome of a Bernoulli experiment.

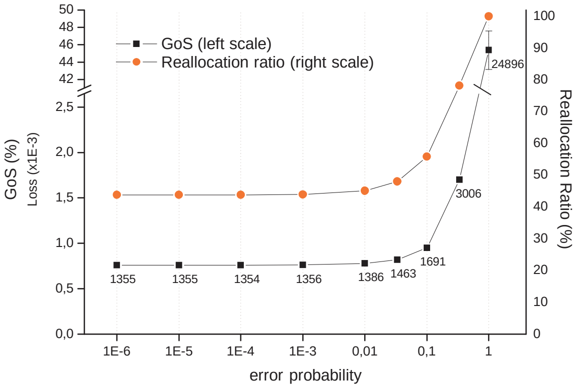

Fig. 4. The reallocation scheme is relatively insensitive to the error probability P of an incorrect reallocation decision.

Fig. 4 presents the GoS, the revenue loss and the reallocation success ratio for several values of the error probability P. Confidence intervals are omitted since they are sufficiently small to be indistinguishable from the plot itself, except for the case $P = 1$. The figure shows that the reallocation scheme is relatively insensitive to the value of the error probability. Even when one in three reallocation decisions is wrong ($P = 0.333$), the network is still better off than without bandwidth reallocation. When nearly all of the reallocation decisions are wrong ($P \sim 1$) the network's performance collapses.

Thus a small probability (in magnitude equal to the order of the packet loss probability in a current generation IP network) of executing incorrect bandwidth reallocation decisions has little impact on the efficiency of the bandwidth reallocation scheme. The scheme would appear to possess an intrinsic capability for self-correcting wrong decisions by reversing incorrect bandwidth transfers upon the next reallocation attempt.

C. *Errors in traffic estimates*

The manager of route r needs to know the connection arrival rate λ_r, the signalling rate η_r, the connection holding time μ_r and the number of connections in service c_r in order to calculate the bandwidth prices. In our simulation experiments, we assume that the values of these parameters are known and we use them in our calculations. In a real implementation they would have to be estimated on-line and errors may arise in their estimation.

In this section we investigate the impact of errors in the estimated value of the connection arrival rate on the efficacy of the reallocation scheme. Let $\lambda'_r = \lambda_r(1 + \delta_r)$ denote an estimate of λ_r where δ_r is a random variable sampled from a uniform distribution in the range $[-X, X]$ where $X < 1$. An value of λ'_r is computed by sampling a new value of δ_r each time the manager of route r computes the buying/selling price of bandwidth.

Fig. 5 presents the GoS and the reallocation success ratio for several values of the arrival rate error parameter X. Confidence intervals are omitted since they are sufficiently small to be indistinguishable from the plot itself. The figure shows that the reallocation scheme is relatively insensitive to reasonable errors in the estimated values of the connection arrival rates. This result is in agreement with [2] where it is shown that it is possible to obtain a robust routing that guarantees a nearly optimal utilization with a fairly limited knowledge of the applicable traffic demands.

V. BANDWIDTH REALLOCATION AMONG SERVICE CLASSES

In this section we investigate another aspect of the robustness of the bandwidth reallocation scheme where the scheme

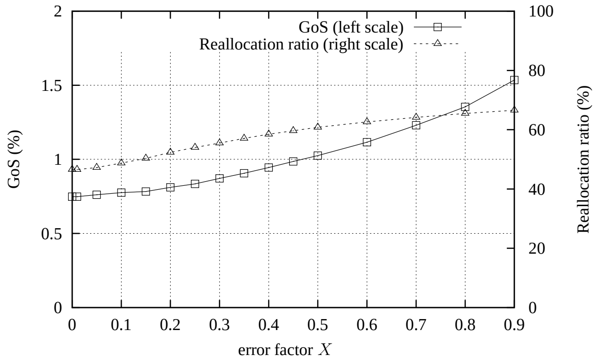

Fig. 5. The reallocation scheme is relatively insensitive to errors in the estimation of the connection arrival rates.

is used to protect the GoS of one traffic *class* against the bandwidth demands of another traffic class.

The calls offered to the network are partitioned into two classes referred to as class 0 and class 1. Two routes connect each OD-pair, one route for class 0 traffic and the other route for class 1 traffic. We consider two forms of bandwidth reallocation, namely *intra-class* and *inter-class* reallocation. In intra-class reallocation, bandwidth reallocation takes place between a transit route and its constituent direct routes, and no reallocation takes place between the class 0 and class 1 constituent direct routes. In this case, the class 0 and class 1 virtual networks are managed independently.

In inter-class reallocation, bandwidth reallocation takes place between a transit route and its constituent direct routes, and in addition reallocation takes place between the constituent direct routes of the different classes as illustrated in Fig. 6. Bandwidth reallocation between the constituent direct routes of the different classes is referred to as *cross-trading*.

The arrival rates $\lambda_r(t)$ of class 1 calls are subjected to a perturbation as described below. The arrival rates λ_r of class 0 calls are left unchanged. Perturbing the arrival rates of class 1 calls only allows us to compare the efficiency of intra-classbandwidth reallocation in the presence or absence of traffic variations. In addition, it will be shown that inter-class trading improves the network's overall GoS, although the GoS of the unperturbed class 0 may suffer when its allocated bandwidth is acquired by the perturbed class 1 in order to better accommodate its varying traffic.

We assume that the arrival rate $\lambda_r(t)$ of class 1 calls to route r is subject to a time-dependent sinusoidal variation. The network model presented in Fig. 2 is used to investigate the effect of inter- and intra-class reallocation. Each simulation terminates after some 15×10^6 calls per class have completed and two sinusoidal cycles have elapsed. Each simulation is independently replicated 50 times. Results as quoted refer to the mean of all replications and confidence intervals are based on a confidence level of 95%. For computational efficiency, the approximate pricing scheme in [5] is used instead of the pricing scheme in [4].

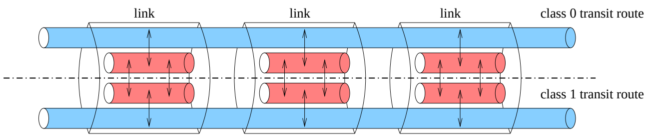

Fig. 6. Inter-class trading: transit routes trade with each other via their respective constituent direct routes.

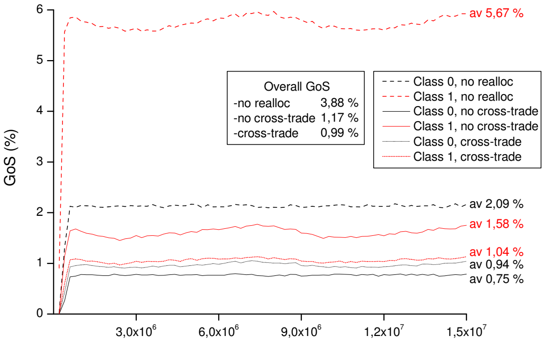

Fig. 7. The results of inter-class trading: the GoS for the unperturbed class 0 and the sinusoidally perturbed (random phases) class 1.

Fig. 7 presents the GoS with intra-class trading. Bandwidth reallocation decreases the GoS of the unperturbed class 0 from 2.09% to 0.75%. The GoS of class 1, which is subject to a sinusoidal traffic perturbation with random phase, is reduced from 5.67% to 1.58%. This demonstrates that the reallocation scheme can provide effective bandwidth assignment despite the traffic perturbation.

Fig. 7 also presents the GoS when inter-class (cross-trading) bandwidth reallocation is enabled. At the expense of a small deterioration in the GoS of the non-perturbed class 0 from 0.75% to 0.94%, the GoS of the perturbed class 1 improves from 1.58% to 1.04%.

Although inter-class trading has a beneficial impact on network performance, class 0 is now indirectly exposed to the traffic fluctuations of class 1 since class 1 can acquire bandwidth from the constituent direct routes of its non-perturbed

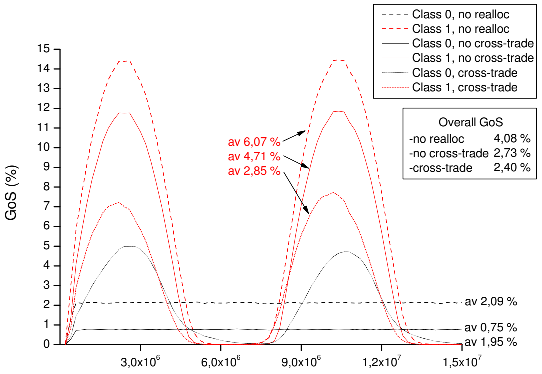

Fig. 8. The results of inter-class trading: the GoS for the unperturbed class 0 and the sinusoidally perturbed (constant phases) class 1.

counterpart. Thus only the perturbed traffic of class 1 strictly benefits.

However, the overall gain from enabling inter-class trading can be redistributed between the two classs. If the buying and selling prices of class 0's constituent direct routes are scaled by a factor of 1.2, then the GoS of class 1 is approximately equal to 0.91 (which is close to the performance without inter-class trading), while the GoS of class 0 improves to 0.72%. Thus it is likely (although not proven) that by choosing a scaling factor between 1.0 and 1.2, arbitrary distributions of the gain resulting from inter-class trading between class 0 and class 1 are possible.

The impact of inter-class trading is further examined in Figs. 8. As before, the network carries two traffic classs. Class 0 is not perturbed and class 1 is subject to a synchronised sinusoidal perturbation. Fig. 8 presents the GoS of both call classes when no bandwidth reallocation takes place. The figure shows that that the GoS of both classs is substantially improved when intra-class bandwidth reallocation takes place. The figure shows that the GoS of both classs improves further when inter-class bandwidth reallocation takes place. However, the inter-class trading allows the indirect reallocation of bandwidth between the two classes and the sinusoidal perturbation of class 1 is communicated to the unperturbed class 0.

VI. CONCLUSIONS

This paper presents a distributed scheme for on-line bandwidth reallocation in a connection-oriented transport network. A bandwidth manager is assigned to each route. The managers are autonomous, acting without centralised control from a system coordinator and behave entirely according to local rules. The managers are aware of local resource demands and bandwidth prices. The managers reallocate bandwidth among themselves in order to maintain the GoS of their routes.

We present a series of simulation experiments of the bandwidth reallocation scheme when applied to a model of a 29-node 45-link network model. The simulation studies reveal that bandwidth reallocation can provide robust and efficient bandwidth provisioning. In addition to improvements in GoS and revenue loss, the reallocation scheme also provides a robust and flexible method for coping with both random and systematic traffic fluctuations. The reallocation scheme can also be used during failure conditions to move bandwidth rapidly from failed routes to recovery routes [7].

REFERENCES

[1] J. Abate and W. Whitt, “Numerical inversion of Laplace transforms of probability distributions”, *ORSA Journal on Computing*, vol. 7, pp. 36–43, 1995.

[2] D. Applegate and E. Cohen. “Making intra-domain routing robust to changing and uncertain traffic demands: understanding fundamental tradeoffs”. Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications, Karlsruhe, Germany (2003) pp 313–324.

[3] Å. Arvidsson, B.A. Chiera, A.E. Krzesinski and P.G. Taylor, “A Distributed Scheme for Value-Based Bandwidth Re-Configuration”, submitted, 2006. `www.cs.sun.ac.za/~aek1/COE/dowloads/four_authors.pdf`

[4] B.A. Chiera and P.G. Taylor, “What is a Unit of Capacity Worth?”, *Probability in the Engineering and Informational Sciences*, vol. 16 no. 4, pp. 513–522, 2002.

[5] B.A. Chiera, A.E. Krzesinski and P.G. Taylor, “Some Properties of the Capacity Value Function”, *SIAM Journal on Applied Mathematics*, vol. 65 no. 4, pp 1407–1419, 2005.

[6] T. Chihara, *An Introduction to Orthogonal Polynomials*. Gordon and Breach, Science Publishers Inc., 1978.

[7] Göbel, A.E. Krzesinski and D. Stapelberg. “A Distributed Scheme for Responsive Network Engineering”. To appear in Proceedings ICC-2007 (July 2007) Glasgow, UK.

[8] S. Kandula, D. Katabi, B. Davie and A. Charny, “Walking the Tightrope: Responsive yet Stable Traffic Engineering”, *ACM SIGCOMM Computer Communication Review*, vol. 35, issue 4, 2005, pp. 253–264.

[9] E. Rosen, A. Viswanathan and R. Callon. Multiprotocol Label Switching Architecture. RFC 3031 (Jan 2001).

[10] T. Anderson, R. Mahajan, N. Spring and D. Wetherall. `www.cs.washington.edu/research/networking/rocketfuel`

Johannes Göbel is a PhD student in the Department of Informatics at the University of Hamburg, Germany. **Dieter Stapelberg** is an MSc student in the Department of Mathematical Sciences at the University of Stellenbosch, South Africa. **Anthony Krzesinski** is a Professor of Computer Science at the University of Stellenbosch, South Africa.