Dynamic Resource Allocation in Communication Networks^{*}

Antonio Capone¹, Jocelyne Elias², Fabio Martignon³, and Guy Pujolle²

¹ Department of Electronics and Information, Politecnico di Milano
`capone@elet.polimi.it`

² University of Paris 6, LIP6 Laboratory,
8 rue du Capitaine Scott, 75015, Paris, France
`jocelyne.elias@lip6.fr`, `guy.pujolle@lip6.fr`

³ Department of Management and Information Technology,
University of Bergamo, Italy
`fabio.martignon@unibg.it`

Abstract. Efficient dynamic resource provisioning algorithms are necessary to the development and automation of Quality of Service (QoS) networks. The main goal of these algorithms is to offer services that satisfy the QoS requirements of individual users while guaranteeing at the same time an efficient utilization of network resources. In this paper we introduce a new service model that provides quantitative per-flow bandwidth guarantees, where users subscribe for a guaranteed rate; moreover, the network periodically individuates unused bandwidth and proposes short-term contracts where extra-bandwidth is allocated and guaranteed exclusively to users who can exploit it to transmit at a rate higher than their subscribed rate. To implement this service model we propose a dynamic provisioning architecture for intra-domain Quality of Service networks. We develop an efficient bandwidth allocation algorithm that takes explicitly into account traffic statistics to increase the users' benefit and the network revenue simultaneously. We demonstrate through simulation in realistic network scenarios that the proposed dynamic provisioning model is superior to static provisioning in providing resource allocation both in terms of total accepted load and network revenue.

Keywords: Dynamic Bandwidth Allocation, Autonomic Networks, Service Model.

1 Introduction

Efficient dynamic resource provisioning mechanisms are necessary to the development and automation of Quality of Service networks. In telecommunication networks, resource allocation is performed mainly in a static way, on time scales

^{*} This work is partially supported by the National Council for Scientific Research in Lebanon.

on the order of hours to months. However, statically provisioned network resources can become insufficient or considerably under-utilized if traffic statistics change significantly [1].

Therefore, a key challenge for the deployment of Quality of Service networks is the development of solutions that can dynamically track traffic statistics and allocate network resources efficiently, satisfying the QoS requirements of users while aiming at maximizing, at the same time, resource utilization and network revenue. Recently, dynamic bandwidth allocation has attracted research interest and many algorithms have been proposed in the literature [1, 2, 3, 4, 5]. These approaches and related works are discussed in Section 2.

Since dynamic provisioning algorithms are complementary to admission control algorithms [1], in our work we assume the existence of admission control algorithms at the edge of the network that cooperate with our proposed bandwidth allocation algorithm operating inside the network.

In this paper we propose a new service model that provides quantitative per-flow bandwidth guarantees, where users subscribe for a guaranteed transmission rate. Moreover, the network periodically individuates unused bandwidth and proposes short-term contracts where extra-bandwidth is allocated and guaranteed exclusively to users who can better exploit it to transmit at a rate higher than their subscribed rate.

To implement this service model we propose a distributed provisioning architecture composed by core and edge routers; core routers monitor bandwidth availability and periodically report this information to ingress routers using signalling messages like those defined in [2]. Moreover, if persistent congestion is detected, core routers notify immediately ingress routers.

Ingress routers perform a dynamic tracking of the effective number of active connections, as proposed in [6], as well as of their actual sending rate. Based on such information and that communicated by core routers, ingress routers allocate network resources dynamically and efficiently using a modified version of the max-min fair allocation algorithm proposed in [7]. Such allocation is performed taking into account users' profile and willingness to acquire extra-bandwidth based on their bandwidth utility function. The allocation is then enforced by traffic conditioners that perform traffic policing and shaping.

We evaluate by simulation the performance of our proposed bandwidth allocation algorithm in realistic network scenarios. Numeric results show that our architecture allows to achieve better performance than statically provisioned networks both in terms of accepted load and network revenue.

In summary, this paper makes the following contributions: the definition of a new service model and the proposition of a distributed architecture that performs dynamic bandwidth allocation to maximize users utility and network revenue.

The paper is structured as follows: Section 2 discusses related work; Section 3 presents our proposed service model and provisioning architecture; Section 4 describes the proposed dynamic bandwidth allocation algorithm; Section 5 discusses simulation results that show the efficiency of our dynamic resource

allocation algorithm compared to a static allocation technique. Finally, Section 6 concludes this work.

2 Related Work

The problem of bandwidth allocation in telecommunication networks has been addressed in many recent works. In [7] a max-min fair allocation algorithm is proposed to allocate bandwidth equally among all connections bottlenecked at the same link. In our work we extend the max-min fair allocation algorithm proposed in [7] to perform a periodical allocation of unused bandwidth to users who expect more than their subscribed rate.

Dynamic bandwidth provisioning in Quality of Service networks has recently attracted a lot of research attention due to its potential to achieve efficient resource utilization while providing the required quality of service to network users [1, 2, 3, 4].

In [1], the authors propose a dynamic core provisioning architecture for differentiated services IP networks. The core provisioning architecture consists of a set of dynamic node and core provisioning algorithms for interior nodes and core networks, respectively. The node provisioning algorithm adopts a self-adaptive mechanism to adjust service weights of weighted fair queuing schedulers at core routers while the core provisioning algorithm reduces edge bandwidth immediately after receiving a Congestion-Alarm signal from a node provisioning module and provides periodic bandwidth re-alignment to establish a modified max-min bandwidth allocation to traffic aggregates.

The work discussed in [1] has similar objectives to our dynamic bandwidth allocation algorithm. However, their service model differs from our proposed model and traffic statistics are not taken into account in the allocation procedure. Moreover, in our work we suggest a distributed architecture implementation, while in these papers only a centralized scheme is considered.

A policy-based architecture is presented in [3], where a measurement-based approach is proposed for dynamic Quality of Service adaptation in DiffServ networks. The proposed architecture is composed of one Policy Decision Point (PDP), a set of Policy Enforcement Points that are installed in ingress routers and bandwidth monitors implemented in core routers. When monitors detect significant changes in available bandwidth they inform the PDP which changes dynamically the policies on in-profile and out-of-profile input traffics based on the current state of the network estimated using the information collected by the monitors. However, this scheme, while achieving dynamic QoS adaptation for multimedia applications, does not take into account the users utility function and their eventual willingness to be charged for transmitting out of profile traffic, thus increasing network revenue.

In [2], a generic pricing structure is presented to characterize the pricing schemes currently used in the Internet, and a dynamic, congestion-sensitive pricing algorithm is introduced to provide an incentive for multimedia applications to adapt their sending rates according to network conditions. As in [2], we take

into account users bandwidth utility functions to evaluate our proposed allocation algorithm based on the increased network revenue that is achieved. However, the authors consider a different service model than that proposed in our work and focus mainly on the issue of dynamic pricing to perform rate adaptation based on network conditions.

The idea of measuring dynamically the effective number of active connections as well as their actual sending rate is a well accepted technique [4, 6]. In [4], the authors propose an active resource management approach (ARM) for differentiated services environment. The basic concept behind ARM is that by effectively knowing when a client is sending packets and how much of its allocated bandwidth is being used at any given time, the unused bandwidth can be reallocated without loss of service. This concept is in line with our proposed bandwidth allocation algorithm. Differently from our work, however, ARM does not guarantee to the user a minimum subscribed bandwidth throughout the contract duration since unused bandwidth is sent to a pool of available bandwidth and it can be used to admit new connections in the network, in spite of those already admitted.

3 Service Model and Dynamic Provisioning Architecture

We first introduce our proposed service model, then we present a distributed provisioning architecture which implements such service model by performing the dynamic bandwidth allocation algorithm described in Section 4; finally, we present the signalling messages used to assure the interaction between network elements.

3.1 Service Model

We propose a service model that, first, provides a quantitative bandwidth guarantee to users and then exploits the unused bandwidth individuated periodically in the network to propose short-term guaranteed extra-bandwidth. In this process, different weights can be assigned to network users to allocate extra-bandwidth with different priorities; such weights can be set statically offline, based on the service contract proposed to the user, or can be adapted on-line based, for example, on the user bandwidth utility function.

Our proposed service model is therefore characterized by:

- a quantitative bandwidth guarantee, expressed through the specification of user's subscribed rate;
- short term guaranteed extra-bandwidth: the network is monitored on-line to individuate unused bandwidth that is allocated with guarantee, during the update interval, to users who can exploit it to transmit extra-traffic;
- a weight that expresses the user's priority in the assignment of extra-bandwidth;
- a bandwidth utility function, $U(x)$, that describes the user's preference for an allocation of x bandwidth units. In line with [8] we consider the utility function as part of the service model. Without loss of generality, we do not consider the pricing component of a bandwidth utility function.

3.2 Architecture and Control Messaging

To implement our service model we assume a distributed architecture constituted by core and edge routers, as shown in Fig.1; traffic monitors are installed on ingress and core routers to perform on-line measurements on the incoming traffic flows and network capacity utilization, respectively.

Core routers exchange messages with ingress routers to report the link utilization or to notify a congestion situation. Each ingress router collects the measurements performed by traffic monitors and exchanges periodically update messages with all other ingress routers to report the current incoming traffic statistics. Moreover, a dynamic bandwidth allocation algorithm is implemented in all ingress routers: it takes into account the traffic statistics gathered at ingress routers and the network information reported by core routers to allocate network resources dynamically and efficiently.

The messages exchanged between network routers, illustrated with arrows in Fig.1, are similar to the control messages that have been proposed in [1] to report persistent congestion or resource availability.

Fig. 1. The proposed distributed architecture that supports dynamic bandwidth allocation

4 Dynamic Bandwidth Allocation Algorithm

We propose a novel dynamic provisioning algorithm that allocates network capacity efficiently based on traffic statistics measured on-line. Bandwidth allocation is performed by ingress routers periodically and is enforced using traffic conditioners. We denote the interval between two successive allocations performed by the algorithm as the *update interval*, whose duration is T_u seconds. Moreover, core routers monitor link utilization, and if congestion on some links is detected, bandwidth re-allocation is immediately invoked to solve this situation.

In the following we present in details the bandwidth allocation algorithm, that proceeds in two steps: in the first step, bandwidth is allocated to all active connections trying to match their near-term traffic requirements that are predicted based on statistics collected by ingress routers. In step two, spare bandwidth as well as bandwidth left unused by idle and active connections is individuated on each link. Such available extra-bandwidth is allocated with guarantee during the current update interval exclusively to connections that can take advantage of it since they are already fully exploiting their subscribed rate.

To illustrate the allocation algorithm, let us model the network as a directed graph $G = (N, L)$ where nodes represent routers and directed arcs represent links. Each link $l \in L$ has associated the capacity C_l . A set of K connections is offered to the network. Each connection is represented by the notation (s_k, d_k, sr_k) , for $k = 1, \dots, K$, where s_k , d_k and sr_k represent the connections source node, destination node and the subscribed rate, respectively; furthermore, we assume that each connection has associated r_min_k , which represents the minimum bandwidth the application requires. Let a_k^l be the routing matrix: $a_k^l = 1$ if connection k is routed on link l , $a_k^l = 0$ otherwise. We assume that a communication between a user pair is established by creating a session involving a path that remains fixed throughout the user pair conversation duration. The session path choice method (i.e., the routing algorithm) is not considered in this paper.

At the beginning of the $n - th$ update interval, each ingress router computes the transmission rate, b_k^{n-1} , averaged over the last T_u seconds, for all connections $k \in K$ that access the network through it. This information is then sent to all other ingress routers using control messages as described in the previous Section, so that all ingress routers can share the same information about current traffic statistics and perform simultaneously the same allocation procedure.

The amount of bandwidth allocated to each source k during the $n - th$ update interval, r_k^n , is determined using the two-steps approach described in the following:

- First step: Connections having $b_k^{n-1} < r_min_k$ are considered *idle*; all other active connections are further classified as *greedy* if they used a fraction greater than γ of their subscribed rate sr_k (i.e. if $b_k^{n-1} > \gamma \cdot sr_k$), otherwise they are classified as *non - greedy*. In our implementation we set $\gamma = 0.9$.

Let us denote by K_i , K_{ng} and K_g the sets of idle, non-greedy and greedy connections, respectively.

Idle connections are assigned their minimum required transmission rate, i.e. $r_k^n = r_min_k, \forall k \in K_i$.

Non-greedy connections are assigned a bandwidth that can accommodate traffic growth in the current update interval while, at the same time, save unused bandwidth that can be re-allocated to other users. Several techniques have been proposed in the literature to predict the near-term transmission rate of a connection based on past traffic measurements. In this work we only consider the last measured value, b_k^{n-1} , and we propose the following simple bandwidth allocation: $r_k^n = \min\{2 \cdot b_k^{n-1}, sr_k\}, \forall k \in K_{ng}$. In this

regard we are currently studying more efficient traffic predictors that could allow improved bandwidth allocation.

Greedy connections are assigned in this step their subscribed rate sr_k , and they also take part to the allocation of extra-bandwidth performed in step two, since they are already exploiting all their subscribed rate.

- Second step: after having performed the allocations described in step one, the algorithm individuates on each link l the residual bandwidth R_l , i.e. the spare bandwidth as well as the bandwidth left unused by idle and non-greedy connections. R_l is hence given by the following expression:

$$R_l = C_l - \left(\sum_{k \in K_i \cup K_{ng}} r_k^n \cdot a_k^l + \sum_{k \in K_g} sr_k \cdot a_k^l \right), \forall l \in L \quad (1)$$

where the first summation represents the total bandwidth allocated in step one to idle and non-greedy connections, while the second summation represents the bandwidth allocated to greedy connections.

Such extra-bandwidth is distributed exclusively to greedy connections using the algorithm detailed in Table 1, which is an extension of the allocation algorithm proposed in [7]. This algorithm takes as input the set K_g of greedy connections, the link set L with the residual capacity on each link l , R_l , and the routing matrix a_k^l , and produces as output the amount of extra-bandwidth f_k^n , $k \in K_g$ that is assigned to each greedy connection during the n -th update interval, so that finally $r_k^n = sr_k + f_k^n$, $\forall k \in K_g$.

To take into account users weights it is sufficient to substitute n_l in Table 1 with w_l , which is defined as the sum of the weights of all greedy connections that are routed on link l .

Table 1. Pseudo-code specification of the bandwidth allocation algorithm

(1) initiate all $f_k^n = 0$, $\forall k \in K_g$
(2) remove from the link set L all links $l \in L$ that have a number of connections crossing them n_l equal to 0
(3) for every link $l \in L$, calculate $F_l = R_l/n_l$
(4) identify the link α that minimizes F_α i.e. $\alpha \mid F_\alpha = \min_k (F_k)$
(5) set $f_k^n = F_\alpha$, $\forall k \in K_\alpha$, where $K_\alpha \subseteq K_g$ is the set of greedy connections that cross link α
(6) for every link l , update the residual capacity and the number of crossing greedy connections as follows:
$R_l = R_l - \sum_{k \in K_\alpha} f_k^n \cdot a_k^l$ $n_l = n_l - \sum_{k \in K_\alpha} a_k^l$
(7) remove from set L link α and those that have $n_l = 0$
(8) if L is empty, then stop; else go to Step (3)

It should be clarified that our algorithm can temporarily present some limitations in bandwidth allocation, since the bandwidth allocated to a user can at most double from an update interval to the successive one. This could affect the performance of users that experience steep increases in their transmission rate. In Section 5 we evaluate numerically this effect showing at the same time how it is counterbalanced by increased network revenue in all the considered network scenarios under several traffic load conditions.

5 Numeric Results

In this Section we compare the performance, measured by the average accepted load and network extra-revenue versus the total load offered to the network, of the proposed dynamic bandwidth allocation algorithm with a static provisioning strategy, referring to different network scenarios to cover a wide range of possible environments. Static provisioning allocates to each source k its subscribed rate sr_k .

We are interested in measuring the following performance metrics: the average accepted load and network extra-revenue. The average accepted load is obtained averaging the total load accepted in the network over all the bandwidth update intervals.

We define, in line with [2], the average network extra-revenue as the total charge paid to the network for all the extra-bandwidth utilization, averaged over all the bandwidth update intervals. In this computation we consider only network extra-revenue generated by greedy users that are assigned extra-bandwidth by our proposed dynamic allocation algorithm. Furthermore we assume, in line with [5], that the utilities are additive so that the aggregate utility of rate allocation is given by the sum of the utilities perceived by all network users.

Using the notation introduced in the previous section, the average network extra-revenue can be obtained averaging over all the update intervals n the quantity:

$$\sum_{k \in K_g} U(b_k^n) - U(sr_k) \quad (2)$$

In the first scenario we gauge the effectiveness of the proposed traffic-based bandwidth allocation algorithm. We consider, in line with [1, 2], the scenario illustrated in Figure 2, that consists of a single-bottleneck with 12 source-destination pairs. All links are full-duplex and have a propagation delay of 1 ms. The capacities of the links are given in the figure.

We use 12 Exponential On-Off traffic sources; the average On time is set to 200 s, and the average Off time is varied in the 0 to 150 s range to simulate different traffic load conditions while at the same time varying the percentage of bandwidth left unused by every connection. Six sources have a peak rate of 40 kb/s and a subscribed rate of 100 kb/s while the remaining sources have a peak rate of 1 Mb/s and a subscribed rate of 300 kb/s; the minimum bandwidth required by each source, r_{min_k} , is equal to 10 kb/s. The algorithm updating interval, T_u , is set equal to 20 s. We assume, for simplicity, that all users have the

Fig. 2. Network topology with a single bottleneck

same weight w_k and the same utility function proposed in [8], $U(x) = 1 - e^{-\frac{x^2}{x+h}}$, that models the perceived utility of real-time elastic traffic for an allocation of x bandwidth units. The parameter h setting is the same as in [8].

Note that a realistic characterization of network applications is outside the scope of this paper. The specification of the utility function allows us exclusively to gauge the extra network revenue that can derive from the deployment of our proposed bandwidth allocation algorithm.

Figures 3(a) and 3(b) show, respectively, the average total load accepted in the network and the corresponding total extra-revenue as a function of the average total load offered to the network. It can be observed that our dynamic provisioning algorithm is very efficient in resource allocation compared to a static provisioning algorithm for all values of the offered load, providing improvements up to 60% in the total accepted traffic.

The maximum network extra-revenue is achieved when the average Off time of exponential sources is equal to 100 s, corresponding to an offered load approximately equal to 4 Mb/s. In this situation, in fact, the average number of idle connections (i.e. 4) is sufficiently high to exalt our dynamic allocation algorithm

Fig. 3. Average total accepted load (a) and network extra-revenue (b) versus the average load offered to the network of Fig. 2

Fig. 4. Average total accepted load (a) and network extra-revenue (b) versus the rate scaling factor α in the network of Fig. 2

that reallocates unused bandwidth to active users who can take advantage of it, sending extra-traffic and generating network revenue. With lower Off time values (i.e. with higher offered loads) the total extra-revenue slightly decreases as less connections are idle, in average, and consequently less bandwidth is available for re-allocation.

To investigate the impact on the performance of the update interval duration, we have considered, in the same scenario, different values for T_u , i.e. 40 s and 60 s. We found that the average increase in the total accepted load, expressed as a percentage of the traffic admitted in the static allocation case, was of 32% for $T_u = 40$ s and 21% for $T_u = 60$ s, while for $T_u = 20$ s it was 47% (see Fig. 3(a)). These results allow to gauge the trade-off between performance improvement and overhead resulting from a more frequent execution of the allocation algorithm.

In the same scenario of Figure 2 we then fixed the average Off time of Exponential sources to 100 s while maintaining the average On time equal to 200 s, and we varied the peak rate of all sources scaling them by a factor α , with

Fig. 5. Network topology with a larger number of links

Table 2. Peak rate, subscribed rate and path for the connections in the network scenario of Figure 5

Connection	Peak Rate (kb/s)	Subscribed Rate (kb/s)	Path
1	100	250	1-3-4
2	100	250	3-4-6
3	100	250	3-4-5
4	1000	500	1-2-4-6
5	1000	500	2-4-5
6	1000	1000	1-3-5

$0.25 \leq \alpha \leq 1.5$. Figures 4(a) and 4(b) show the total accepted load and the total extra-revenue in this scenario.

At very low load the static provisioning technique achieves slightly higher performance than dynamic provisioning. This is due to the fact that in this situation static provisioning is in effect sufficient to accommodate all incoming traffic; on the other hand, dynamic provisioning needs some time (in the worst case up to T_u seconds) to track the transition of sources from the idle to the active state. For all other traffic loads the advantage of the proposed dynamic bandwidth allocation algorithm is evident both in terms of accepted load and network extra-revenue.

A more realistic scenario is shown in Fig. 5. It comprises 6 nodes and 8 bidirectional links, all having a capacity equal to 2 Mb/s and propagation delay of 1 ms. In this topology, 6 Exponential On-Off traffic sources are considered, and their source and destination nodes are indicated in the Figure. Table 2 reports for all the connections the peak rate, the subscribed rate and the path of the connection. All other parameters are set as in the previous scenarios. Note that, with such paths choice, various connections compete for network capacity with different connections on different links.

Fig. 6. Average total accepted load (a) and network extra-revenue (b) versus the average load offered to the network of Fig. 5

Also in this scenario the dynamic allocation algorithm outperforms static allocation, as shown in Figures 6(a) and 6(b), thus proving the benefit of the proposed scheme. These results verify that our allocation algorithm allows service providers to increase network capacity utilization and consequently network extra-revenue with respect to static provisioning techniques.

6 Conclusion

In this paper we proposed a novel service model where users subscribe for guaranteed transmission rates, and the network periodically individuates unused bandwidth that is re-allocated and guaranteed with short-term contracts to users who can better exploit it. We described a distributed dynamic resource provisioning architecture for quality of service networks. We developed an efficient bandwidth allocation algorithm that takes explicitly into account traffic statistics to increase the users perceived utility and the network extra-revenue.

Simulations results measured in realistic network scenarios show that our allocation algorithm allows to increase both resource utilization and network revenue with respect to static provisioning techniques.

References

1. A. T. Campbell and R. R.-F. Liao. Dynamic Core Provisioning for Quantitative Differentiated Services. *IEEE/ACM Transactions on Networking*, pages 429–442, vol. 12, no. 3, June 2004.
2. H. Schulzrinne and X. Wang. Incentive-Compatible Adaptation of Internet Real-Time Multimedia. *IEEE Journal on Selected Areas in Communications*, pages 417–436, vol. 23, no. 2, February 2005.
3. T. Ahmed, R. Boutaba, and A. Mehaoua. A Measurement-Based Approach for Dynamic QoS Adaptation in DiffServ Network. *Journal of Computer Communications, Special issue on End-to-End Quality of Service Differentiation, Elsevier Science*, 2004.
4. M. Mahajan, M. Parasharand, and A. Ramanathan. Active Resource Management for the Differentiated Services Environment. *International Journal of Network Management*, pages 149–165, vol. 14, no. 3, May 2004.
5. F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, pages 33–37, vol. 8, 1997.
6. J. Aweya, M. Ouellette, and D. Y. Montuno. Design and stability analysis of a rate control algorithm using the Routh-Hurwitz stability criterion. *IEEE/ACM Transactions on Networking*, pages 719–732, vol. 12, no. 4, August 2004.
7. D. Bertsekas and R. Gallager. *Data Networks, 2nd Edition*. Prentice-Hall, 1992.
8. L. Breslau and S. Shenker. Best-Effort versus Reservations: A Simple Comparative Analysis. In *Proc. ACM SIGCOMM*, pages 3–16, September 1998.