Impact of Network Topology on the Performance of Budget Based Network Admission Control Methods

Michael Menth[1], Stefan Kopf[1], and Joachim Charzinski[2]

[1] Department of Distributed Systems, Institute of Computer Science, University of Würzburg, Am Hubland, 97074 Würzburg, Germany
{menth,kopf}@informatik.uni-wuerzburg.de
[2] Siemens AG, Germany
joachim.charzinski@siemens.com

Abstract. Budget based network admission control (NAC) mechanisms can be categorized into four basically distinct approaches. Since they have different complexity and efficiency, we compare their resource utilization in different networking scenarios. Our results show that the network size, the connectivity, and the internal structure of the network have a significant impact on the resource efficiency. Some NAC approaches can achieve a very high utilization if the offered load is large enough while the performance of others is limited by the network topology. This study does not focus on specific protocols because the presented NAC schemes classify most existing resource management schemes. It is intended to optimize the NAC design for future QoS networks.

1 Introduction

The next generation of the Internet is expected to fully integrate all kinds of data and media communications. In contrast to today's telephone network, data connections have variable bitrates and the management of the individual nodes should be simpler. And in contrast to today's Internet, real-time multimedia applications expect mechanisms for increased Quality of Service (QoS). This implies that future networks need a limitation of traffic load [1] to meet the packet loss and delay requirements. This function is called admission control (AC). High quality transmission is guaranteed at the expense of control, management effort, and blocked reservation requests in overload situations. To realize a low border-to-border (b2b) flow blocking probability in transit networks, the networks are provided with sufficient transport capacities which causes costs for the network provider. Therefore, AC mechanisms should be efficient but still simple. For reasons of robustness, they should not induce information states inside the network.

Link admission control (LAC) limits the transported traffic on a single link to prevent violations of its QoS requirements. In contrast, *network* admission

This work was funded by the Bundesministerium für Bildung und Forschung of the Federal Republic of Germany (Förderkennzeichen 01AK045) and Siemens AG, Munich. The authors alone are responsible for the content of the paper.

control (NAC) is required when data are transported over several hops through a network instead just over a single link. This may be done by applying LAC on a link-by-link basis but this implies AC states in the core. However, it is desirable to control the load inside the network only at the border routers by performing AC based on resource budgets that are prereserved for certain traffic aggregates. In this work we identify four different NAC methods that reveal different resource utilization and that categorize most of today's implemented and investigated NAC approaches.

The paper is structured as follows. Section 2 gives an overview of four basic budget based NAC categories. Section 3 shows their achievable resource utilization in different networking scenarios and analyzes their performance. Section 4 summarizes this work and gives an outlook on further research.

2 Methods for Network Admission Control (NAC)

In this section we introduce four different budget based NAC concepts. A NAC instance records the demand of all admitted active flows $\mathcal{F}_{admitted}$. When a new flow arrives, AC checks whether its effective bandwidth together with the demand of already established flows fits within a capacity budget. If so, the flow is accepted, otherwise it is rejected. For the sake of a simple description, we take only peak rate allocation for flows into account. However, all mechanisms can be combined with more efficient LAC methods like the use of effective bandwidth or measurement based AC [2–4].

2.1 Link Budget Based Network Admission Control (LB NAC)

The link-by-link NAC is probably most intuitive. The capacity[3] $l.c$ of each link l in the network is managed by a single link budget (LB) $LB(l)$ with size $LB(l).c$ that may be administered, e.g. at the ingress router of that link or in a centralized database. A new flow $f^{v,w}_{new}$ with ingress router v[4], egress router w, and bitrate $f^{v,w}_{new}.c$ must pass the AC procedure for the LBs of all links that are traversed in the network by $f^{v,w}_{new}$ (cf. Figure 1(a)). The NAC procedure is successful if the following inequality holds

$$\forall l \in \mathcal{E} : l.u(v,w) > 0 \ : \quad f^{v,w}_{new}.c \cdot l.u(v,w) + \sum_{f^{x,y} \in \mathcal{F}_{admitted}(l)} f^{x,y}.c \cdot l.u(x,y) \leq LB(l).c. \tag{1}$$

[3] We borrow parts of our notation from the object-oriented programming style: $x.y$ denotes a property y of an object x. We prefer $x.y$ to the conventional y_x since this is hard to read if the name of x is complex.

[4] A networking scenario $\mathcal{N} = (\mathcal{V}, \mathcal{E}, u)$ is given by a set of border routers $\mathcal{V}$ and set of links $\mathcal{E}$. The b2b traffic aggregate with ingress router v and egress router w is denoted by $g(v,w)$. The function $l.u(v,w)$ with $v,w \in \mathcal{V}$ and $l \in \mathcal{E}$ reflects the routing and it is able to cover both single- and multi-path routing by indicating the percentage of the traffic rate $g(v,w).c$ using link l.

There are many systems and protocols working according to that principle. The connection AC in ATM [5] and the Integrated Services [6, 7] architecture in IP technology adopt it in pure form. Other protocols reveal the same behavior although the mechanism is not implemented as an explicit LB NAC. A bandwidth broker [8–10] administers the budgets in a central entity which represents a single point of failure but behaves in a similar way. The stateless-core approaches [11–13] are able to avoid states in the core at the expense of measurements or increased response time.

With this approach, core routers need to hold AC states per flow which is problematic with respect to scalability and robustness. If network resilience is required, flows are deviated in case of a partial network outage and their AC states must be quickly restored. This would entail a tremendous technical overhead in real-time for large systems. The following three NAC methods manage the network capacity in a distributed way, i.e. all budgets related to a flow can be consulted at its ingress or its egress border router. In a failure scenario, only fast local rerouting of the traffic is required if sufficient backup capacities are available.

2.2 Ingress Budget and Egress Budget Based Network Admission Control (IB/EB NAC)

The IB/EB NAC defines for every ingress node $v \in \mathcal{V}$ an ingress budget (IB) $IB(v)$ and for every egress node $w \in \mathcal{V}$ an egress budget (EB) $EB(w)$ that must not be exceeded. A new flow $f_{new}^{v,w}$ must pass the AC procedure for $IB(v)$ and $EB(w)$ and it is only admitted if the requests to both budgets are successful (cf. Figure 1(b)). Hence, the following inequalities must hold

$$f_{new}^{v,w}.c + \sum_{f \in \mathcal{F}_{admitted}^{ingress}(v)} f.c \leq IB(v).c \quad \text{and} \tag{2}$$

$$f_{new}^{v,w}.c + \sum_{f \in \mathcal{F}_{admitted}^{egress}(w)} f.c \leq EB(w).c \tag{3}$$

Flows are admitted at the ingress and the egress irrespective of their egress or ingress routers. The mere IB NAC, which originates from the DiffServ context [14–16], admits traffic only at the ingress border router and only Equation (2) must be met for the AC procedure. Capacity managed by an *IB* or *EB* can be used in a very flexible manner. However, the network must be able to carry all – also pathological – traffic patterns that are acceptable by the IBs and EBs with the required QoS. Therefore, enough capacity must be allocated such that also very unlikely – but admissible – scenarios with a strongly skewed traffic matrix can be supported.

2.3 B2B Budget Based Network Admission Control (BBB NAC)

A b2b budget (BBB) $BBB(v, w)$ manages the capacity for all flows between v and w, i.e. it defines a virtual tunnel in case of single-path routing. Hence, the

(a) LB NAC.

(b) IB/EB NAC.

(c) BBB NAC.

(d) ILB/ELB NAC.

Fig. 1. Budget based network admission control (NAC) methods.

BBB NAC takes both the ingress and the egress border router of a flow $f^{v,w}$ into account for the AC procedure. A new flow $f^{v,w}_{new}$ passes only the AC procedure for $BBB(v,w)$ (cf. Figure 1(c)). It is admitted if this request is successful, i.e. if the following inequality holds

$$f^{v,w}_{new}.c + \sum_{f \in \mathcal{F}_{admitted}(v,w)} f.c \leq BBB(v,w).c. \tag{4}$$

The $BBB(v,w)$ may be controlled at the ingress router v or at the egress router w, i.e. the BBB NAC avoids states in the core, too. Tunnels may also be used hierarchically [17, 18]. The tunnel capacity may be signaled using explicit reservation states in the network [19, 20], only in logical entities like bandwidth brokers [9], or it may be assigned by a central entity [21]. The capacity $BBB(v,w).c$ of a tunnel is dedicated to one specific b2b aggregate $g(v,w)$ and can not be used for other traffic with different source or destination. Hence, there is no flexibility for resource utilization but pathological traffic patterns are excluded. The BBB NAC is often implemented in a more flexible manner, such that the size of the BBBs can be rearranged [22, 23]. The same can be done for the other NACs, too.

2.4 Ingress Link Budget and Egress Link Budget Based Network Admission Control (ILB/ELB NAC)

The ILB/ELB NAC defines ingress link budgets (ILBs) $ILB(l,v)$ and egress link budgets (ELBs) $ELB(l,w)$ to manage the capacity of each $l \in \mathcal{E}$. They are administered by border routers v and w, i.e. the link capacity is partitioned among $|\mathcal{V}|-1$ border routers. In case of single-path IP routing, the links $\{l : \mathit{ILB}(l,v) > 0\}$ constitute a source tree and the links $\{l : \mathit{ELB}(l,w) > 0\}$ form a sink tree (cf. Figure 1(d)). A new flow $f^{v,w}_{new}$ must pass the AC procedure for the $\mathit{ILB}(.,v)$ and $\mathit{ELB}(.,w)$ of all links that are traversed in the network by $f^{v,w}_{new}$ (cf. Figure 1(d)). The NAC procedure will be successful if the following inequalities are fulfilled

$$\forall l \in \mathcal{E} : l.u(v,w) > 0 :$$
$$f^{v,w}_{new}.c \cdot l.u(v,w) + \sum_{f^{v,y} \in \mathcal{F}^{l,v,ingress}_{admitted}} f^{v,y}.c \cdot l.u(v,y) \leq \mathit{ILB}(l,v).c \quad \text{and} \tag{5}$$
$$\forall l \in \mathcal{E} : l.u(v,w) > 0 :$$
$$f^{v,w}_{new}.c \cdot l.u(v,w) + \sum_{f^{x,w} \in \mathcal{F}^{l,w,egress}_{admitted}} f^{x,w}.c \cdot l.u(x,w) \leq \mathit{ELB}(l,w).c. \tag{6}$$

There are several significant differences to the BBB NAC. A BBB covers only an aggregate of flows with the same source and destination while the ILBs (ELBs) may cover flows with the same source (destination) but different destinations (sources). Therefore, the ILB/ELB NAC is more flexible than the BBB NAC. With the BBB NAC, only one $\mathit{BBB}(v,w)$ is checked while with ILB/ELB NAC, the number of budgets to be checked is twice the path length of a flow. Like with the IB/EB NAC, there is the option to use only ILBs or ELBs by applying only Equation (5) or Equation (6). The ILB/ELB or ILB NAC can be viewed as local bandwidth brokers at the border routers, disposing over a fraction of the network capacity. These concepts are new and have not yet been implemented by any resource management protocol. The token based distributed NAC resembles the ILB NAC if it works in the responsive mode [13] from a performance point of view. Although the path of the sessions in BGRP [24] matches also a sink tree, BGRP works like the LB NAC on its entities.

3 Performance Comparison of NAC Approaches

In this section the capacity of sample networks is dimensioned to meet a desired blocking probability p_{b2b} in presence of a given traffic matrix. This is done for all NAC methods according to the formulae in [25] to evaluate the sum of the required link capacities and the resulting resource utilization. We take these values as performance measures in our study. Most observations in this work are due to the notion of multiplexing gain or economy of scale. This is the fact that a larger offered load leads to a more efficient provisioning of a resource.

3.1 Influence of the Offered Load

Our performance evaluation framework for NAC methods is based on queuing theory and it is described in detail in [25]. To study the impact of the offered load on the NAC performance, we take the test network depicted in Figure 2. Its topology is based on the UUNET in 1994 [26] where nodes connected by only one or two links to the network were successively removed. To model real-time connections in the Internet, the flows themselves have heterogeneous rate requests. We assume a homogeneous traffic matrix and scale it by the offered b2b load a_{b2b} which is the average number of flows between two border routers. In this investigation, shortest path routing is used. Due to these limitations, we investigate the impact of traffic matrix and routing on the performance of NAC methods in [27].

Figure 3 shows the resource utilization depending on the offered load a_{b2b} for all NAC methods. The LB NAC uses the network resources most efficiently. A budget $LB(l)$ controls a maximum possible amount of traffic on link l and takes most advantage from economy of scale. The ILB/ELB, ILB, and BBB NAC are less efficient because the same offered load $g(v,w).a \cdot l.u(v,w)$ is partitioned among up to $|\mathcal{V}|$ budgets in case of ILB NAC or $|\mathcal{V}| \cdot (|\mathcal{V}|-1)$ different budgets in case of BBB NAC. The reduced traffic load per budget leads to smaller multiplexing gain and requires more overall capacity $l.c$ for the same link. For sufficiently high offered load, the utilization of all these NAC methods tends towards 100%. Some NACs are not able to exclude unlikely traffic patterns which force to allocate high link capacities to an extent that reduces the achievable resource utilization to 30% for the IB/EB NAC and to 10% for the IB NAC. Hence, the IB NAC has the worst performance and our IB/EB NAC achieves a three times larger resource utilization by applying the limitation of the traffic volume in a symmetric way.

Fig. 2. Test network.

Fig. 3. The impact of the offered load on the resource utilization.

3.2 Influence of the Network Topology

The network topology is another factor influencing the NAC performance. The resource efficiency depends on the average node degree, the network size, and on the internal structure of the network.

Construction of Random Networks. The degree $deg(v)$ of a node $v \in \mathcal{V}$ is the number of links connected to this node and the average node degree of a network can be calculated by $deg_{avg} = \frac{2 \cdot |\mathcal{E}|}{|\mathcal{V}|}$. The authors of [28] propose algorithms for the random construction of inter-networks. However, we use our own construction methods (CM) because we consider only a single autonomous system and we want to control the node degree quite rigidly. Since we want to have a decentralized network, we set the maximum node degree to $deg_{max} = deg_{avg} + 1$. The CMs of our random networks respect these constraints and avoid loops and parallels. They start by building a spanning tree network and continue with one of the following options.

- CM0 connects nodes with a largest distance.
- CM1 connects nodes randomly.
- CM2 connects nodes with a shortest distance.

If not mentioned differently, we choose CM1 for our studies, we set the network size to $|\mathcal{V}| = 50$ and the average node degree to $deg_{avg} = 5$. We use a small offered load $a_{b2b} = 10$ to make the difference between the NAC types more visible since the resource utilization of some of them converges for large a_{b2b} to 100%. For each data point we analyzed 10 different random networks to obtain small confidence intervals that are omitted in the figures.

Influence of the Network Size. Figure 4(a) illustrates how the required network capacity and the average path length rise with the network size $|\mathcal{V}|$. The growth is mainly due to our traffic model, i.e. the overall offered load scales quadratically with the number of nodes. The number of links grows only linearly ($|\mathcal{E}| = \frac{|\mathcal{V}| \cdot deg_{avg}}{2}$). Hence, there is a linear growth of the offered load per link below the line, not yet taken into account that the average path length rises as well with increasing network size. Figure 4(b) reveals that only the link budget based NAC methods (LB, ILB, ILB/ELB NAC) can take advantage of traffic concentration caused by an increased number of b2b aggregates and achieves a larger resource utilization. For the sake of clarity, we omitted the curves for the ILB NAC in the figures whose resource efficiency and capacity requirements lie between the ILB/ELB NAC and the BBB NAC. The resource utilization of the BBB NAC remains constant since the traffic load offered to the budget equals the entries in the traffic matrix (a_{b2b}) and does not change. This underlines again the advantage of the ILB and ILB/ELB NAC approaches: although they do not induce states in the core, their performance can benefit from an increased traffic volume due to a larger network size. The performance of the IB NAC is low and decreases with increasing network size. The IB/EB NAC is inefficient and the resource efficiency decreases with the network size but it still outperforms the IB NAC significantly.

(a) (b)

(c) (d)

Fig. 4. The sensitivity of the required network capacity and the resource utilization.

Influence of the Average Node Degree. Figure 4(c) shows that the required network capacity for the LB and the BBB NAC is clearly dominated by the average path length, which directly correlates with the overall traffic volume in the network. The same holds for the ILB and the ILB/ELB NAC.

The required capacity for the IB NAC is independent of the average node degree deg_{avg}. The IP routing tree seen by any source node is a spanning tree consisting of $(|\mathcal{V}| - 1)$ edges if all routers are both core and border routers.

Each of the links of the spanning tree must support the full capacity $IB(v).c$ to avoid congestion in the case that the NAC admits a traffic pattern whose rate is $IB(v).c$ and all the traffic goes over that link.

The IB NAC admits traffic aggregates with a rate of $IB(v).c$ and all of its traffic going over any specific link of the routing tree. This scenario is unlikely but can not be excluded by the IB NAC. Hence, the required network capacity induced by a single ingress node v is $(|\mathcal{V}| - 1) \cdot IB(v).c$ and the required network capacity is $\sum_{v \in \mathcal{V}} (|\mathcal{V}| - 1) \cdot IB(v).c$. Therefore, the required capacity is independent of the topology as long as the number of routers is constant.

The IB/EB NAC restricts pathologic traffic patterns more efficiently than the IB NAC and requires less capacity. However, it is remarkable that the required capacity rises with increasing node degree although the average path length decreases. The following reveals that this is due to the internal structure of the network.

Influence of Hierarchical Structures. Figure 4(d) illustrates that the average path length depends significantly on the average node degree and the CM. The network capacity required for the IB/EB NAC rises with increasing node degree. Although CM2 leads to the longest paths and to most traffic in the network, it requires clearly less capacity than CM0 and CM1. A node degree of 2 yields almost a spanning tree network where $deg_{avg}=2 \cdot \frac{(|\mathcal{V}|-1)}{|\mathcal{V}|} \approx 2$. Since the spanning tree is the base for all CMs, the required capacity for the IB/EB NAC is about the same for all CMs for a node degree of $deg_{avg}=2$.

We analyze these observations. CM0 tries to add as many shortcuts as possible to the initial spanning tree which results in a relatively short path length. Randomly constructed networks lead to approximately the same results. However, CM2 avoids the installation of efficient shortcuts and yields a larger average path length than CM0 and CM1. Therefore, the initial spanning tree structure dominates the CM2 topology and leads to a kind of traffic backbone since many shortest paths in the network use the links of the original spanning tree. Hence, CM2 networks reveal some hierarchical structure.

To explain the reduced capacity requirements for CM2 for the IB/EB NAC, we consider the link l of a router w with node degree $deg(w)=1$. The $IB(v), v \in \mathcal{V} \setminus w$, limit the required capacity for that link to $l.c=\sum_{\{v\in\mathcal{V}:v\neq w\}} IB(v).c$. In addition, the required capacity can be limited by $l.c=EB(w).c$, too. In case of a homogeneous traffic matrix, we have $IB(v).c=EB(w).c$ for all $v, w \in \mathcal{V}$, hence, the required capacity for link l is reduced to $\frac{1}{(|\mathcal{V}|-1)}$ of the capacity required for the IB NAC.

If a link is used by other cross traffic, too, the limitation of the required link capacity by the EBs is not so efficient. An increasing average node degree increases the number of links, it makes most nodes transit nodes for multiple flows by providing shortcuts, thereby reducing the hierarchical structure of the network. Thus, if the node degree rises, the traffic limitation by the combination of the IBs and EBs looses efficiency to a certain extent. And this is more likely to happen with CM0 and CM1 than with CM2.

Finally, the IB/EB NAC can benefit from hierarchical network structures to limit the required capacity. However, its performance is still restricted to low values. This experiment also shows the sensitivity of the performance of the IB/EB NAC to the internal network structure beyond network size and node degree.

4 Conclusion

We distinguished between *link* admission control (LAC) and *network* admission control (NAC). LAC limits the number of flows on a single link to assure their

QoS requirements while NAC limits the number of flows in a network. We presented four basic NAC methods: the link budget (LB) based NAC, the border-to-border (b2b) budget (BBB) based NAC, which consists of virtual tunnels, the ingress and egress budget (IB/EB) based NAC, known from the Differentiated Services context, and the ingress and egress link budget (ILB/ELB) based NAC, which is a new concept. Many research projects implement admission control (AC) schemes that can be classified by these categories.

For each NAC method, we dimensioned the capacity of sample networks to meet a desired blocking probability in presence of a given traffic matrix. The NAC types revealed significantly different resource efficiency which is mainly due to their ability for taking advantage of economy of scale. The LB NAC exhibits the best resource utilization, followed by the ILB/ELB NAC, the ILB NAC, and the BBB NAC. However, they all achieve a resource utilization close to 100% if the offered traffic load is sufficiently high. The IB and IB/EB NAC are less efficient as they achieve a utilization in the order of $\frac{apl}{|\mathcal{V}|}$ where $|\mathcal{V}|$ is the number of border routers in the network and *apl* the average path length. The concepts of ILB NAC, ILB/ELB NAC, and IB/EB NAC are new and they outperform the BBB NAC and the mere IB NAC, respectively. For a fixed b2b offered load (a_{b2b}), the NAC performance depends clearly on the network size, the average node degree, and – in particular for the IB/EB NAC – on the internal structure of the network.

This work presented a first evaluation of NAC methods and the focus was the sensitivity of the NAC performance to the network topology. Currently, we evaluate the effect of skewed traffic matrices and different routing schemes. We work on optimal strategies for the capacity assignment of the budget sizes in presence of limited link capacities [29]. If local network outages occur, the QoS of the traffic can not be maintained unless the traffic is quickly rerouted. This, however, implies backup capacity in the network which raises the question for resource efficiency of NAC methods under resilience requirements [30]. Optimized routing mechanisms can further improve the resource efficiency under

References

1. Shenker, S.: Fundamental Design Issues for the Future Internet. IEEE JSAC **13** (1995) 1176–1188
2. Roberts, J., Mocci, U., Virtamo, J.: Broadband Network Teletraffic - Final Report of Action COST 242. Springer, Berlin, Heidelberg (1996)
3. Breslau, L., Jamin, S., Shenker, S.: Comments on the Performance of Measurement-Based Admission Control Algorithms. In: Infocom. (2000) 1233–1242 ISBN 0-7803-5880-5.
4. Kelly, F.P.: Notes on Effective Bandwidths. Stochastic Networks: Theory and Applications **4** (1996) 141–168
5. The ATM Forum: Traffic Management Specification, Version 4.0. (1996)
6. Braden, B., Clark, D., Shenker, S.: RFC1633: Integrated Services in the Internet Architecture: an Overview. http://www.ietf.org/rfc/rfc1633.txt (1994)

7. Menth, M., Martin, R.: Performance Evaluation of the Extensions for Control Message Retransmissions in RSVP. In: 7^{th} International Workshop on Protocols for High-Speed Networks (PfHSN 2002), Berlin, Germany (2002)
8. Terzis, A., Wang, J., Ogawa, J., Zhang, L.: A Two-Tier Resource Management Model for the Internet. In: Global Internet Symposium'99. (1999)
9. Teitelbaum, B., Hares, S., Dunn, L., Narayan, V., Neilson, R., Reichmeyer, F.: Internet2 QBone: Building a Testbed for Differentiated Services. IEEE Network Magazine (1999)
10. Zhang, Z.L.Z., Duan, Z., Hou, Y.T.: On Scalable Design of Bandwidth Brokers. IEICE Transaction on Communications **E84-B** (2001) 2011–2025
11. Stoica, I., Zhang, H.: Providing Guaranteed Services Without Per Flow Management. Computer Communication Review **29** (1999)
12. Szábó, R., Henk, T., Rexhepi, V., Karagiannis, G.: Resource Management in Differentiated Services (RMD) IP Networks. In: International Conference on Emerging Telecommunications Technologies and Applications (ICETA 2001), Kosice, Slovak Republic (2001)
13. Bhatnagar, S., Nath, B.: Distributed Admission Control to Support Guaranteed Services in Core-Stateless Networks. In: IEEE INFOCOM 2003, San Francisco, USA (2003)
14. Blake, S., Black, D.L., Carlson, M.A., Davies, E., Wang, Z., Weiss, W.: RFC2475: An Architecture for Differentiated Services. ftp://ftp.isi.edu/in-notes/rfc2475.txt (1998)
15. Xiao, X., Ni, L.M.: Internet QoS: A Big Picture. IEEE Network Magazine **13** (1999) 8–18
16. Engel, T., Nikolouzou, E., Ricciato, F., Sampatakos, P.: Analysis of Adaptive Resource Distribution Algorithsm in the Framework of a Dynamic DiffServ IP Network. In: 8^{th} International Conference on Advances in Communications and Control (ComCon8), Crete, Greece (2001)
17. Menth, M., Hauck, N.: A Graph-Theoretical Notation for the Construction of LSP Hierarchies. In: 15^{th} ITC Specialist Seminar, Würzburg, Germany (2002)
18. Kompella, K., Rekhter, Y.: LSP Hierarchy with Generalized MPLS TE. http://www.ietf.org/internet-drafts/draft-ietf-mpls-lsp-hierarchy-08.txt (2002)
19. Baker, F., Iturralde, C., Le Faucheur, F., Davie, B.: RFC3175: Aggregation of RSVP for IPv4 and IPv6 Reservations. http://www.ietf.org/rfc/rfc3175.txt (2001)
20. Awduche, D.O., Berger, L., Gan, D.H., Li, T., Srinivasan, V., Swallow, G.: RFC3209: RSVP-TE: Extensions to RSVP for LSP Tunnels. http://www.ietf.org/rfc/rfc3209.txt (2001)
21. Trimintzios, P., Bauge, T., Pavlou, G., Georgiadis, L., Flegkas, P., Egan, R.: Quality of Service Provisioning for Supporting Premium Services in IP Networks. In: IEEE Globecom 2002, Taipei, Taiwan (2002)
22. Menth, M.: A Scalable Protocol Architecture for End-to-End Signaling and Resource Reservation in IP Networks. In: 17^{th} International Teletraffic Congress, Salvador de Bahia, Brazil (2001) 211–222
23. Fu, H., Knightly, E.: Aggregation and Scalable QoS: A Performance Study. In: Proceedings of IWQoS 2001, Karlsruhe, Germany (2001)
24. Pan, P., Schulzrinne, H.: BGRP: A Tree-Based Aggregation Protocol for Interdomain Reservations. Journal of Communications and Networks **2** (2000) 157–167
25. Menth, M., Kopf, S., Milbrandt, J.: A Performance Evaluation Framework for Network Admission Control Methods. Technical Report, No. 305, University of Würzburg, Institute of Computer Science (2003) http://www-info3.informatik.uni-wuerzburg.de/TR/tr305.pdf.

26. : Network Maps. (http://www.nthelp.com/maps.htm)
27. Menth, M., Kopf, S., Milbrandt, J.: Impact of Traffic Matrix and Routing on the Performance of Network Admission Control Methods. Technical Report, No. 307, University of Würzburg, Institute of Computer Science (2003) http://www-info3.informatik.uni-wuerzburg.de/TR/tr307.pdf.
28. Zegura, E.W., Calvert, K.L., Donahoo, M.J.: A Quantitative Comparison of Graph-Based Models for Internet Topology. IEEE/ACM Transactions on Networking **5** (1997) 770–783
29. Menth, M., Gehrsitz, S., Milbrandt, J.: Fair Assignment of Efficient Network Admission Control Budgets. Technical Report, No. 306, University of Würzburg, Institute of Computer Science (2003) http://www-info3.informatik.uni-wuerzburg.de/TR/tr306.pdf.
30. Menth, M., Kopf, S., Charzinski, J.: Impact of Resilience Requirements on the Performance of Network Admission Control Methods. Technical Report, No. 309, University of Würzburg, Institute of Computer Science (2003) http://www-info3.informatik.uni-wuerzburg.de/TR/tr309.pdf.