Network Admission Control for Fault-Tolerant QoS Provisioning

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, Munich, Germany
`joachim.charzinski@siemens.com`

Abstract. In a connection oriented network layer, admission control (AC) is easily combined with connection state management at each network node. However, after a link or node failure, existing connections are dropped or reservations must be restored on new paths, which requires high signalling effort. In contrast, a connectionless network layer like IP does not deal with connection or resource management at the network nodes. After a failure, connectivity is easily restored by rerouting, affecting higher layer connections only via some packet drops. Thus, a resource management scheme for IP should allow rerouting to cope with failures without affecting reservation states. A *network* admission control (NAC) handles reservations only at dedicated locations, e.g. the borders of a network, not burdening individual routers with admission decisions or reservation states. The NAC architecture enables *resilient* resource reservation, maintaining reservations even after failures and intra-domain rerouting. In this paper, we investigate the efficiency of three different distributed budget management schemes with single and multi-path routing. We show how the admission decision can be designed to be tolerable against failure scenarios by admitting only the amount of traffic that can still be carried after a failure and the corresponding rerouting.

1 Introduction

A next generation Internet is expected to fully integrate all kinds of data and media communications. In contrast to today's telephone network, applications have variable bitrate requirements and the management of the individual nodes should be simpler. And in contrast to today's Internet, broadband real-time applications require a minimum Quality of Service (QoS). This implies that in future networks the traffic load must be limited [1] to meet applications' bit rate and delay requirements. The corresponding function is called admission control (AC). High quality transmission is guaranteed at the expense of blocking reservation requests in overload situations.

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.

Networks are dimensioned such that for a given traffic matrix the blocking probability is small enough not to upset customers while keeping link capacities and thus cost as low as possible. Introducing a QoS reservation architecture does not make much sense unless the QoS is also maintained throughout periods of failures of network elements like links, routers, or router interfaces. Whereas intra-domain routing protocols like OSPF [2] can quickly restore connectivity in an IP network, current resource reservation architectures do not ensure that after rerouting there is sufficient bandwidth available on the new paths for the existing reservations. Unlike with physical layer protection, the corresponding additional capacity can be used for best-effort traffic in normal operation, which makes resilience on the network layer cheaper than on the physical layer. A simple option to exploit this feature in a protection context would be to use MPLS fast rerouting [3], abandoning the flexibility and adaptivity of IP routing.

Network admission control (NAC) schemes allow keeping resource reservation states separate from the routers. In [4] we have identified several fundamentally distinct NAC categories which reveal different resource efficiency. Link-by-link NAC budgets similar to ATM or IntServ [5, 6] may be managed in a centralized database; ingress and/or egress rate budgets may be allocated to border routers like in the DiffServ context [7]; or the network resources may be managed as virtual tunnels [8, 9]. In this paper we show how these NAC schemes can be used to provide resilient resource reservation by preventively including failure scenarios in the budgets. We investigate the efficiency of three different NAC schemes and compare their efficiency under resilience requirements, using the single-path and Equal Cost Multi-Path (ECMP) variants of shortest path routing as in OSPF.

The paper is structured as follows. Section 2 gives an overview of three basic budget based NAC categories. Section 3 explains how suitable budget and link capacities can be dimensioned and how to include resilience requirements in NAC budgets. Section 4 compares the resource efficiency of NAC methods for networks with and without backup capacity as well as for single- and multi-path routing.

2 Methods for Network Admission Control (NAC)

In this section we distinguish between link and network admission control and explain three basically different NAC concepts.

2.1 Link and Network Admission Control

QoS criteria are usually formulated in a probabilistic way, i.e., the packet loss probability and the probability that the transport delay of a packet exceeds a given delay budget must both be lower than certain thresholds. Link admission control (LAC) takes the queuing characteristics of the traffic into account and determines the required bandwidth to carry flows over a single link without QoS violations.

Network admission control (NAC) needs to protect more than one link with one admission decision. This is a distributed problem with various solutions differing in their degree of storage and processing demands, locality and achievable multiplexing gain due to the partitioning of resources into budgets administered in different locations.

Moreover, the solutions have different efficiency, i.e. they require different amounts of network capacity to meet the same border-to-border (b2b) flow blocking probability p_{b2b} which affects the network operator's costs.

NAC and LAC can be combined, i.e. a flow's required capacity $c(f)$ may consist of an effective bandwidth to take burstiness and/or some overbooking in the presence of large traffic aggregates into account. In this investigation, we only focus on the combinatoric NAC problem, i.e. we work on effective bandwidth budgets and blind out the issues of determining the effective bandwidth for individual reservations or potential MBAC based overbooking.

In general, an AC entity records the demand of the admitted flows $\mathcal{F}(b)$ in place related to a budget b. When a new flow arrives, it checks whether its effective bandwidth together with the demand of already established flows fits within the capacity budget. If so, the flow is accepted, otherwise it is rejected. This principle is used in link based admission control, controlling one link, as well as as in NAC where a number of network resources are covered by each budget and at the same time the utilization of one resource is affected by a number of budgets.

2.2 Link Budget Based Network Admission Control (LB NAC)

The link-by-link NAC is probably the most intuitive NAC approach. The capacity $c(l)$ of each link l in the network is managed by a single link budget LB_l (with size $c(LB_l)$) that may be administered, e.g., at the router sending over that link or in a centralized database. A networking scenario $\mathcal{N} = (\mathcal{V}, \mathcal{E}, u)$ is given by a set of 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 set of all b2b traffic aggregates is $\mathcal{G}$. The function $u_l(g_{v,w})$ indicates the percentage of the traffic rate $c(g_{v,w})$ using link l. It is able to reflect both single- and multi-path routing. A new flow $f^{new}_{v,w}$ with ingress router v, egress router w, and bitrate $c(f^{new}_{v,w})$ must pass the AC procedure for the LBs of all links that are traversed in the network by $f^{new}_{v,w}$ (cf. Fig. 1(a)). The NAC procedure will be successful if the following inequality holds

$$\forall l \in \mathcal{E} : u_l(g_{v,w}) > 0 : c(f^{new}_{v,w}) \cdot u_l(g_{v,w}) + \sum_{f_{x,y} \in \mathcal{F}(LB_l)} c(f_{x,y}) \cdot u_l(g_{x,y}) \leq c(LB_l). \qquad (1)$$

There are many systems and protocols working according to that principle. The connection AC in ATM [5] and the Integrated Services [6] architecture proposed for IP adopt it in pure form and induce per flow reservation states in the core. Other architectures reveal the same behavior although the mechanism is not implemented as an explicit LB NAC. A bandwidth broker [10, 8, 11] administers the budgets in a central database. The stateless core approaches [12–14] avoid reservation states in the core at the expense of measurements or increased response time. Reservation states in the core, measurements, or increased response times are a drawback if network resilience is required. The following three basic 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 traffic rerouting is required and the QoS is maintained if sufficient backup capacity is available.

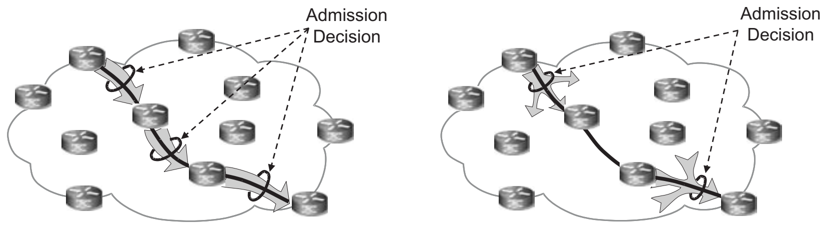

(a) LB NAC. (b) IB/EB NAC.

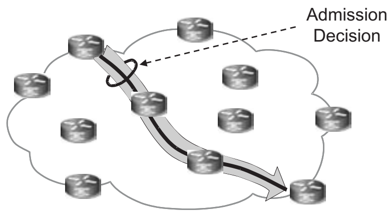

(c) BBB NAC.

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

2.3 Ingress 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_v and for every egress node $w \in \mathcal{V}$ an egress budget EB_w that must not be exceeded. A new flow $f_{v,w}^{new}$ must pass the AC procedure for IB_v and EB_w and it is only admitted if both requests are successful (cf. Fig. 1(b)). Hence, the following inequalities must hold

$$c(f_{v,w}^{new}) + \sum_{f \in \mathcal{F}(IB_v)} c(f) \leq c(IB_v) \quad \text{and} \quad c(f_{v,w}^{new}) + \sum_{f \in \mathcal{F}(EB_w)} c(f) \leq c(EB_w) \qquad (2)$$

Flows are admitted at the ingress irrespective of their egress router and at their egress router irrespective of their ingress routers, i.e. both AC decisions are decoupled. This entails that the 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 – combinations of traffic patterns that are admissible by the IBs and EBs with the required QoS. Hence, sufficient capacity must be allocated or the IBs and EBs must be set small enough.

If we leave the EBs aside, we get the simple IB NAC, so only the left part of Eq. (2) is checked for the AC procedure. This idea fits within the DiffServ context [15, 7] where traffic is admitted only at the ingress routers without looking at the destination address of the flows. The QoS should be guaranteed by a sufficiently low utilization of the network resources by high quality traffic.

2.4 B2B Budget Based Network Admission Control (BBB NAC)

The BBB NAC is able to exclude pathological traffic patterns by taking both the ingress and the egress border router of a flow $f_{v,w}$ into account for the AC decision, i.e. a b2b budget $BBB_{v,w}$ manages the capacity of a virtual tunnel between v and w. This tunnel can consist of multiple b2b paths if multi-path routing is used. Fig. 1(c) illustrates that a new flow $f_{v,w}^{new}$ passes only a single AC procedure for $BBB_{v,w}$. It is admitted if the following inequality holds

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

The BBB NAC can also avoid states inside the network because the $BBB_{v,w}$ may be controlled at the ingress or egress router. The capacity of a tunnel is bound by the BBB to one specific b2b aggregate and can not be used for other traffic with different source or destination. Hence, there is no flexibility for resource utilization. Therefore, the concept is often realized in a more flexible manner, such that the size of the BBBs can be rearranged [16, 17]. Tunnels may also be used hierarchically [18]. The tunnel capacity may be signaled using explicit reservation states in the network [9, 19], only in logical entities like bandwidth brokers [8], or it may be assigned by a central entity [20].

3 Capacity Dimensioning for Budgets and Links

AC guarantees QoS for admitted flows at the expense of flow blocking if the budget capacity is exhausted. Since this applies to all budgets mentioned before, we abstract from special budgets to a general one denoted by b. To keep the blocking probability small, the capacity $c(b)$ of a budget b must be dimensioned large enough. First, we consider budget dimensioning in general. Then, we explain how NAC specific budgets and link capacities are calculated. Finally, we define a performance measure for the comparison of NAC methods and show how to include resilience in NAC budgets.

3.1 Capacity Dimensioning

We review a general approach for capacity dimensioning and derive the required blocking probabilities.

Capacity Dimensioning for a Single Budget Capacity dimensioning is a function calculating the required bandwidth for given traffic characteristics and a desired blocking probability. The specific implementation of that function depends on the underlying traffic model. We assume Poisson arrivals of resource requests and a generally distributed holding time. Although typical Internet traffic has different characteristics on the packet level [21], the Poisson model is more realistic for the resource request level of end-user driven real-time applications. In addition, we are rather interested in a basic performance comparison of the NAC methods than in the capacity dimensioning for a specific network service with known traffic profiles. The offered load a is the mean number of active flows, if no flow blocking occured. In a multi-service world, the request profile is multi-rate, so we take n_r different request types r_i, $0 \leq i < n_r$ with a bitrate $c(r_i)$. Given an offered load a, the respective request type specific offered load is $a(r_i) = p_a(r_i) \cdot a$. In our studies, we assume a simplified multimedia real-time communication scenario with $n_r = 3$, $c(r_0) = 64$ kbit/s, $c(r_1) = 256$ kbit/s, and $c(r_2) = 2048$ kbit/s, and a mean bitrate of $E[C] = \sum_{0 \leq i < n_r} c(r_i) \cdot p_a(r_i) = 256$ kbit/s. The recursive solution by Kaufman and Roberts [22] allows for the computation of request type specific blocking probabilities $p_b(r_i)$ if a certain capacity c is provided. We use Eq. (4) to relate the blocking probability p_b to the traffic volume instead to the number of flows:

$$p_b = 1 - \Big[\sum_{0 \leq i < n_r} (1 - p_b(r_i)) \cdot c(r_i) \cdot p_a(r_i) \Big] / E[C]. \tag{4}$$

An adaptation of the Kaufman and Roberts algorithm yields the required capacity c for a desired blocking probability p_b. After all, we can compute the required budget capacity $c(b)$ if the offered load $a(b)$ and the desired budget blocking probability $p_b(b)$ are given.

From B2B Blocking Probabilities to Budget Blocking Probabilities Budget sizes are dimensioned for a desired budget blocking probability $p_b(b)$. The set $\mathcal{B}_g$ consists of all budgets whose capacity needs to be checked if a flow of the traffic aggregate g asks for admission. The b2b blocking probability associated with this aggregate g is then

$$p_{b2b}(g) = 1 - \Pi_{b \in \mathcal{B}_g} (1 - p_b(b)). \tag{5}$$

under the assumption that flow blocking at different budgets is independent. Since flow blocking at different budgets tends to be positively correlated, the computation of $p_{b2b}(g)$ according to Eq. (5) is rather conservative.

In [4] we have proposed three different methods for setting the budget blocking probabilities $p_b(b)$ to achieve a desired b2b flow blocking probability p_{b2b}. They have hardly any effect on the NAC performance, therefore, we stick with the simple approach that all $p_b(b)$ are equal for all budgets $b \in \mathcal{B}_g$. We denote by $m(b)$ the maximum number of budgets to be checked for any flow controlled by b. Then the required $p_b(b)$ is determined by

$$p_b(b) \leq 1 - \sqrt[m(b)]{1 - p_{b2b}} \tag{6}$$

3.2 Resource Allocation for Budget Based NAC Methods

We denote the offered load for a b2b aggregate by $g_{v,w}$ by $a(g_{v,w})$. The resulting matrix $A_{\mathcal{G}} = \big(a(g_{v,w})\big)_{v,w \in \mathcal{V}}$ is the traffic matrix. In contrast, the current requested rate of an aggregate is $c(g_{v,w})$ and the matrix $C_{\mathcal{G}} = \big(c(g_{v,w})\big)_{v,w \in \mathcal{V}}$ describes an instantaneous traffic pattern. For a possible traffic pattern $C_{\mathcal{G}} \in {\mathbb{R}_0^+}^{|\mathcal{V}|^2}$ the following formulae hold

$$\forall v, w \in \mathcal{V} : c(g_{v,w}) \geq 0 \quad \text{and} \quad \forall v \in \mathcal{V} : c(g_{v,v}) = 0. \tag{7}$$

If NAC is applied in the network, each traffic pattern $C_{\mathcal{G}}$ satisfies the constraints defined by the NAC budgets. These constraints lead to linear equations, too, serving as side conditions for the calculation of the worst case scenario on each link $l \in \mathcal{E}$ by the following rate maximization:

$$c(l) \geq \max_{C_\mathcal{G} \in \mathbb{R}_0^{+|\mathcal{V}|^2}} \sum_{g \in \mathcal{G}} c(g) \cdot u_l(g). \tag{8}$$

This determines the minimum required capacity $c(l)$ of link l. Since the aggregate rates have real values, the maximization can be performed by the Simplex algorithm [23] in polynomial time. However, for some NAC methods there are more efficient solutions that we will point out in the following.

LB NAC The LB NAC requires that a transit flow needs to check a budget LB_l for every link l of its path for admission, hence, the maximum number of passed NAC budgets is

$$m(LB_l) = max_{\{g \in \mathcal{G}: u_l(g) > 0\}} len_{path}^{max}(g, l) \tag{9}$$

whereby $len_{path}^{max}(g, l)$ is the maximum length of a path containing l used by g. As the budget LB_l covers all flows traversing link l, its expected offered load is

$$a(\mathit{LB}_l) = \sum_{g \in \mathcal{G}} a(g) \cdot u_l(g). \tag{10}$$

According to Eq. (1)

$$\forall l \in \mathcal{E} : \sum_{g \in \mathcal{G}} c(g) \cdot u_l(g) \leq c(\mathit{LB}_l) \tag{11}$$

must be fulfilled, so the minimum capacity $c(l)$ of link l is constrained by

$$c(l) \geq c(\mathit{LB}_l). \tag{12}$$

IB/EB NAC With the IB/EB NAC, a flow is admitted by checking both the ingress and the egress budget. Thus, we get $m(\mathit{IB}_v) = m(\mathit{EB}_w) = 2$. The IB/EB NAC subsumes all flows with the same ingress router v under IB_v and all flows with the same egress router w under EB_w. The offered load of the respective budgets is

$$a(\mathit{IB}_v) = \sum_{w \in \mathcal{V}} a(g_{v,w}) \quad \text{and} \quad a(\mathit{EB}_w) = \sum_{v \in \mathcal{V}} a(g_{v,w}). \tag{13}$$

Here we use the inequalities from Eq. (2) as side conditions in Simplex method for the computation of the capacity $c(l)$:

$$\forall v \in \mathcal{V} : \sum_{w \in \mathcal{V}} c(g_{v,w}) \leq c(\mathit{IB}_v) \quad \text{and} \quad \forall w \in \mathcal{V} : \sum_{v \in \mathcal{V}} c(g_{v,w}) \leq c(\mathit{EB}_w). \tag{14}$$

In case of the mere IB NAC, $m(\mathit{IB}_v) = 1$ holds. The IBs are computed in the same way like above, however, there is a computational shortcut to the Simplex method for the calculation of the required link capacity $c(l)$:

$$c(l) \geq \sum_{v \in \mathcal{V}} c(\mathit{IB}_v) \cdot \sum_{w \in \mathcal{V}} u_l(g_{v,w}) \tag{15}$$

BBB NAC With the BBB NAC, only one budget is checked, therefore, $m(BBB_{v,w}) = 1$. The BBB NAC subsumes under $BBB_{v,w}$ all flows with ingress router v and egress router w. The offered load for $BBB_{v,w}$ is simply

$$a(BBB_{v,w}) = a(g_{v,w}). \tag{16}$$

Since Eq. (3) is checked for admission

$$\forall v, w \in \mathcal{V} : c(g_{v,w}) \leq c(BBB_{v,w}) \tag{17}$$

must be fulfilled and the minimum capacity $c(l)$ of link l is constrained by

$$c(l) \geq \sum_{v,w \in \mathcal{V}} c(BBB_{v,w}) \cdot u_l(g_{v,w}) \tag{18}$$

3.3 Performance Measure for NAC Comparison

We compute the required link capacities for all NAC methods according to the equations above. The required network capacity $c(\mathcal{N})$ is the sum of all link capacities in the network. The overall transmitted traffic rate $\hat{c}(\mathcal{N})$ is the sum of the offered load of all b2b aggregates g weighted by their average path lengths $len^{avg}_{path}(g)$, their acceptance probability $(1 - p_{b2b})$, and the mean request rate $E[C]$. We can neglect the fact that requests with a larger rate have a higher blocking probability due to the construction in Eq. (4).

$$c(\mathcal{N}) = \sum_{l \in \mathcal{E}} c(l) \tag{19}$$

$$\hat{c}(\mathcal{N}) = (1 - p_{b2b}) \cdot E[C] \cdot \sum_{\{g \in \mathcal{G}\}} a(g) \cdot len^{avg}_{path}(g) \tag{20}$$

$$\rho(\mathcal{N}) = \frac{\hat{c}(\mathcal{N})}{c(\mathcal{N})}. \tag{21}$$

The overall resource utilization $\rho(\mathcal{N})$ is the fraction of the transmitted traffic rate and the overall network capacity. We use it in the next section as the performance measure for the comparison of NAC methods.

3.4 Resilience Requirements

Even if network resources have been properly assigned to guarantee the QoS of all flows, a local outage in a network can lead to severe QoS problems as rerouting may lead to congestion on other links. Therefore, sufficient capacity must be provisioned beforehand to carry all planned traffic even in the case of outages. Secondly, NAC must limit the admitted resource requests so that the spare capacity required for resilience remains unallocated in normal operation.

Appropriate dimensioning for possible outage scenarios is required which takes care of the rerouted traffic in that case. To that aim, the set $\mathcal{S}$ of protected failure scenarios

$\mathcal{P}$ must be known. Each $\mathcal{P} \in \mathcal{S}$ reflects a set of failed network elements $\mathcal{V}^F_{\mathcal{P}} \subseteq \mathcal{V}$ and $\mathcal{E}^F_{\mathcal{P}} \subseteq \mathcal{E}$, i.e. the set of working routers $\mathcal{V}^W_{\mathcal{P}} \subseteq \mathcal{V}$ and the set of working links $\mathcal{E}^W_{\mathcal{P}} \subseteq \mathcal{E}$ are different from $\mathcal{V}$ and $\mathcal{E}$ which yields a new routing function $u_{\mathcal{P}}$.After all, we have a new networking scenario $\mathcal{N}_{\mathcal{P}}$ for every protected failure scenario $\mathcal{P} \in \mathcal{S}$. We denote $\mathcal{P}$ with $\mathcal{V}^F_{\mathcal{P}} = \emptyset$ and $\mathcal{E}^F_{\mathcal{P}} = \emptyset$ by $\mathcal{P}^*$ and define that it is always contained in $\mathcal{S}$ to facilitate the handling of the normal operation mode in the following. Each link $l \in \mathcal{E}$ must be provided with sufficient capacity to carry the premium traffic in all $\mathcal{P} \in \mathcal{S}$. Hence, the required link capacity is

$$c(l) \geq \max_{\mathcal{P} \in \mathcal{S}} c_{\mathcal{P}}(l). \tag{22}$$

As outlined before, the NAC limits the traffic in the networks by Eq. (1) which leads to the inequalities in Eq. (11), Eq. (14) and Eq. (17) that can be used in a linear program to evaluate the required link capacities. In an outage scenario $\mathcal{P}$, the routing function $u_l(g_{v,w})$ becomes $u_l^{\mathcal{P}}(g_{v,w})$ which must be respected in the traffic maximization step in Eq. (8). As long as the budgets are not changed, the side conditions are still based on the old routing function $u_l^{\mathcal{P}^*}(g_{v,w})$. Due to this change, the shortcut for the calculation of the link capacities for the LB NAC in Eq. (12) does not work anymore and the time consuming Simplex method must be applied.

4 NAC Performance Under Resilience Requirements

We investigate the performance of each NAC method analytically using the above equations, with and without resilience requirements, and with single-path (SP) and multi-path routing (MP) for which we choose shortest single-path routing and shortest equal cost multi-path routing based on a hop count metric. We take SP and MP routing as the routing mechanisms in normal operation mode and use their convergence as reroute mechanism. Therefore, the routing function $u_{\mathcal{P}}$ in a failure scenario $\mathcal{P}$ equals the conventional SP or MP routing in the resulting networking scenario. The number of failure scenarios with n link failures is $\binom{|\mathcal{E}|}{n}$. The more links fail, the less likely is that scenario and the more expensive is its protection. Therefore, we restrict our numerical studies to all single bi-directional link failure scenarios, although any relevant failure scenario could principally be included in the budgets.

We study the NAC performance in the COST 239 network (cf. Fig. 2, [24]) since it allows many shortest equal cost multi-paths to illustrate the influence of MP routing. Our performance measure is the average resource utilization $\rho(\mathcal{N})$. It is limited by three factors: (1) the amount of overdimensioning required to ensure bandwidth availability for the given traffic demand at a given blocking probability, (2) the spare capacity provisioned for the case of failures and (3) the amount of overdimensioning required to accomodate all the combinations of flows that may be admitted by the independent NAC instances.

In the following, we illustrate first the concept of economy of scale on a single link relating to issue (1). Then, we investigate the influence of the routing scheme and the resilience requirements on the resource utilization (issue (2)) depending on the NAC method (issue (3)), and finally, we compare these methods in the different scenarios.

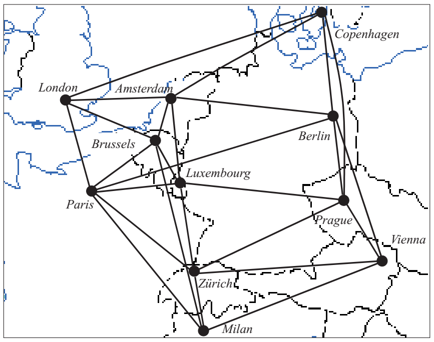

Fig. 2. The topology of the COST239 network.

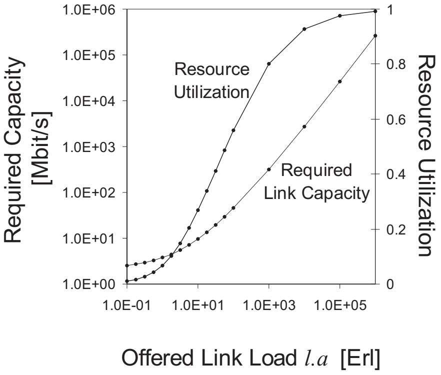

Fig. 3. The impact of offered load on the required link capacity and resource utilization on a single link under link admission control.

4.1 Economy of Scale Illustrated on a Single Link

Economy of scale or multiplexing gain is the key for understanding the performance behavior of any admission control approach and can be best illustrated on a single link. In [4] we have shown that the b2b blocking probability has a minor impact on the required capacity and the resource utilization compared to the offered load. In all our studies we set $p_{b2b} = 10^{-3}$. Fig. 3 shows that the required link capacity and the resource utilization depend heavily on the offered link load $a(l)$. The resource utilization increases drastically up to an offered load of $a(l) = 1000$ Erlang where the economy of scale is fully exploited. Then the required link capacity rises almost linearly with the offered link load. The performance depends also on the network topology, on the routing, and on the traffic matrix which has been studied in [25, 26].

4.2 Impact of Resilience Requirements and Routing on the NAC Methods

BBB NAC Fig. 4(a) shows the resource utilization for the BBB NAC. The average offered load $a(g_{v,w})$ of all b2b aggregates $g_{v,w}$ is given by our system parameter a_{b2b}. Since the BBBs cover exactly that traffic, the performance of the BBB NAC without resilience requirements equals exactly the single link scenario discussed above. The routing does not influence the offered load $a(BBB_{v,w}) = a_{b2b}$ of a budget (cf. Eq. (16)) and the resulting required capacities $c(BBB_{v,w})$ add up to the link capacities (cf. Eq. (18)). Therefore, the overall required network capacity $c(\mathcal{N})$ for the BBB NAC is the same regardless of the routing as long as packets are forwarded on a shortest path.

With resilience requirements only 60% and 68% resource utilization can be achieved in the limit for SP and MP routing, respectively. The reciprocal value $\frac{1}{0.6} \approx 1.67$ is the average degree of overdimensioning required for the survivability in outage scenarios. This corresponds to 67% additional backup capacity. Hence, clearly less than twice the capacity is required to achieve 100% resilience for all outage scenarios because the backup capacity is shared by different flows in different link failure scenarios. MP routing reduces this value even further to less than 50% ($\frac{1}{0.68} \approx 1.47$) since the rerouted traffic is distributed equally to more links which need in turn less backup capacity for resilience purposes. This observation is very important and applies to the performance of other NAC methods under resilience requirements, too.

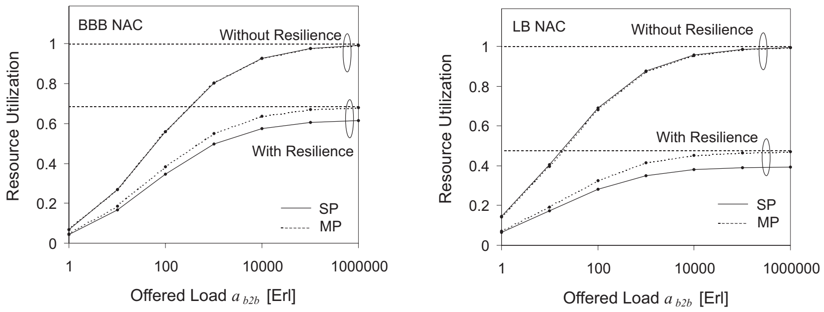

(a) BBB NAC. (b) LB NAC.

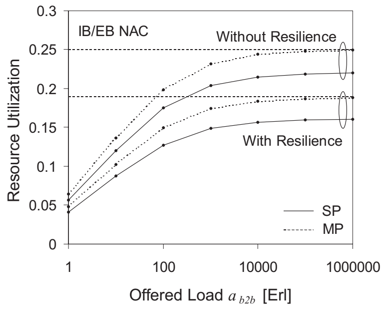

(c) IB/EB NAC.

Fig. 4. Resource utilization in the COST Network for different NAC methods, routing schemes and resilience requirements.

LB NAC Fig. 4(b) illustrates the resource utilization of the LB NAC. Again, the LB NAC performance hardly depends on the routing scheme in the non-resilient case because resource efficiency depends only on the traffic concentration on the links. The routing options SP and MP do not affect the resource utilization sufficiently to achieve clearly visible effects. When resilience requirements are included in the budgets, the resource utilization decreases to 40% for SP routing and 48% for MP routing. Although the absolute utilization values are smaller, the backup resource sharing effect observed with BBBs applies here too.

IB/EB NAC Fig. 4(c) illustrates the achievable utilization with the IB/EB NAC. It is at most 22% for SP routing without resilience requirements and 16% with resilience requirements which leads to only 37.5% additional backup capacities. MP routing also improves the NAC performance under resilience requirements due to the reasons given above. It is remarkable that MP routing increasses the performance also for operation without resilience. However, this kind of NAC is still far less efficient than the BBB NAC.

5 Conclusion

In this paper we have distinguished between link admission control (LAC) and network admission control (NAC). We reviewed three NAC categories and showed how to compute their budgets or dimension link capacities. The novelty in this paper is the consideration of link failure scenarios in the admission decision such that rerouted traffic is still carried with the desired QoS.

The measure for performance comparison is the average resource efficiency, indicating the amount of required spare capacity. We tested the performance of each NAC method with and without resilience requirements, and with hop-count based shortest path routing with its single-path (SP) and equal cost multi-path options.

A direct comparison of the NAC methods without resilience requirements and SP routing shows that the LB NAC is most efficient for low and medium size offered load, followed by the BBB NAC. These NAC types achieve a resource utilization close to 100% for sufficiently high offered load. In contrast, the performance of the IB/EB NAC converges to a network specific asymptote between 10% and 20%. Without resilience requirements, LB NAC, BBB NAC and IB/EB NAC are not influenced by the routing scheme whereas the performance of IB/EB NAC is improved by MP routing. Under resilience requirements, the efficient NAC methods achieve a lower resource utilization between 40% and 70%. They have different utilization limits and the order of their efficiency is reversed, i.e. the BBB NAC is more efficient than the LB NAC. Under resilience requirements, all NAC methods profit from multi-path routing. We have observed the same effects in different network topologies with different utilization limits.

Hence, networks resilient against element failures should implement the BBB NAC for two reasons. First, the network has a stateless core and no resource reservation signalling is needed when traffic is rerouted. Second, the BBB NAC requires less backup capacity than any other NAC approach. In addition, the capacity calculation for BBB NAC is easier and its implementation is less complex compared to other NAC methods.

With MP routing, the resource utilization is about 10% larger than with SP routing. This shows that the mechanism for rerouting in failure scenarios holds some optimization potential with regard to the amount of the required backup capacity. Both MPLS or enhanced MP routing schemes may be used to increase the resource utilization and to reduce the required backup capacity in resilient networks.

References

1. Shenker, S.: Fundamental Design Issues for the Future Internet. IEEE JSAC **13** (1995) 1176–1188

2. Moy, J.: RFC2328: OSPF Version 2. ftp://ftp.isi.edu/in-notes/rfc2212.txt (1998)
3. Pan, P., Gan, D.H., Swallow, G., Vasseur, J.P., Cooper, D., Atlas, A., Jork, M.: Fast Reroute Extensions to RSVP-TE for LSP Tunnels. http://www.ietf.org/internet-drafts/draft-ietf-mpls-rsvp-lsp-fastreroute-03.txt (2003)
4. Menth, M., Kopf, S., Milbrandt, J.: A Performance Evaluation Framework for Network Admission Control Methods. In: IEEE Network Operations and Management Symposium (NOMS), Seoul, South Korea (2004)
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. Xiao, X., Ni, L.M.: Internet QoS: A Big Picture. IEEE Network Magazine **13** (1999) 8–18
8. Teitelbaum, B., Hares, S., Dunn, L., Narayan, V., Neilson, R., Reichmeyer, F.: Internet2 QBone: Building a Testbed for Differentiated Services. IEEE Network Magazine (1999)
9. 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)
10. Terzis, A., Wang, J., Ogawa, J., Zhang, L.: A Two-Tier Resource Management Model for the Internet. In: Global Internet Symposium'99. (1999)
11. Zhang, Z.L.Z., Duan, Z., Hou, Y.T.: On Scalable Design of Bandwidth Brokers. IEICE Transaction on Communications **E84-B** (2001) 2011–2025
12. Stoica, I., Zhang, H.: Providing Guaranteed Services Without Per Flow Management. Computer Communication Review **29** (1999)
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. 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)
15. 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)
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 Conf. on Advances in Commun. and Control (ComCon8), Crete, Greece (2001)
17. Fu, H., Knightly, E.: Aggregation and Scalable QoS: A Performance Study. In: Proceedings of IWQoS 2001, Karlsruhe, Germany (2001)
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. 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)
20. 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)
21. Paxson, V., Floyd, S.: Wide-Area Traffic: The Failure of Poisson Modeling. IEEE/ACM Transactions on Networking **3** (1995) 226–244
22. Roberts, J., Mocci, U., Virtamo, J.: Broadband Network Teletraffic - Final Report of Action COST 242. Springer, Berlin, Heidelberg (1996)
23. Stoer, J.: Numerische Mathematik 1. 5th edn. Springer, New York, Berlin, Heidelberg (1989)
24. Batchelor et al., P.: Ultra High Capacity Optical Transmission Networks. Final report of Action COST 239. http://barolo.ita.hsr.ch/cost239/network/ (1999)
25. Menth, M., Kopf, S., Charzinski, J.: Impact of Network Topology on the Performance of Network Admission Control Methods. In: MIPS2003, Naples, Italy (2003) 195 – 206
26. Menth, M., Milbrandt, J., Kopf, S.: Impact of Routing and Traffic Distribution on the Performance of Network Admission Control. In: ISCC2004, Alexandria, Egypt (2004)