Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.

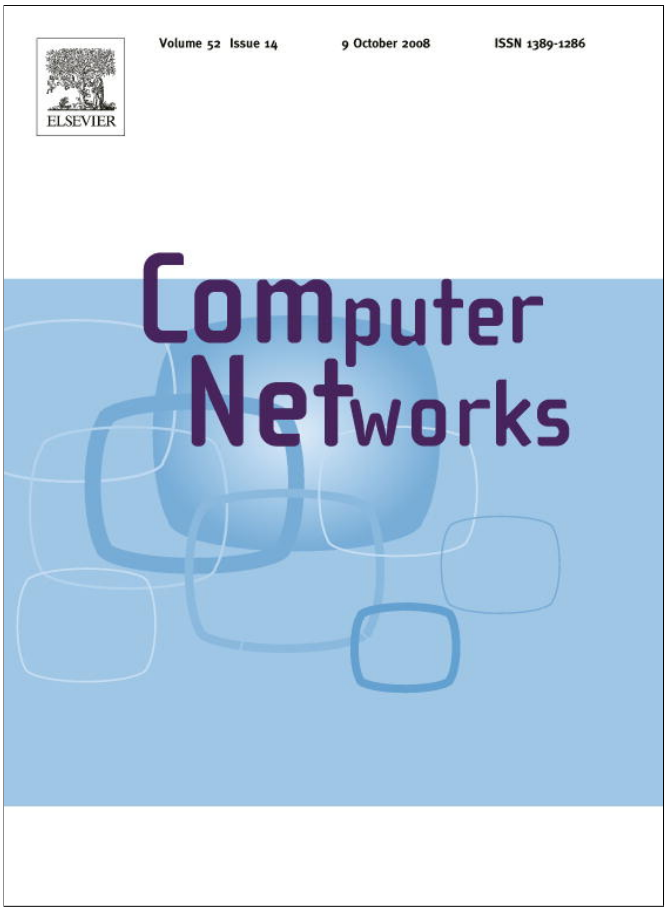

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

©ACM 2008. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in Computer Networks.

Author's personal copy

Computer Networks 52 (2008) 2805–2815

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Resilient network admission control

Michael Menth [a,*], Stefan Kopf [a], Joachim Charzinski [b], Karl Schrodi [b]

[a] University of Würzburg, Institute of Computer Science, Germany
[b] Nokia Siemens Networks, Munich, Germany

ARTICLE INFO

Article history:
Received 24 March 2007
Received in revised form 28 December 2007
Accepted 30 May 2008
Available online 25 June 2008

Responsible Editor: M. Smirnow

Keywords:
Admission control
Resilience
Network dimensioning
QoS

ABSTRACT

Network admission control (NAC) limits the traffic in a network to avoid overload and to assure thereby the quality of service (QoS) for admitted flows. Overload may occur due to exceptional traffic demand, but it is mostly caused by redirected traffic due to link failures. Conventional NAC methods cannot cope with network outages and fail when they are needed most. This paper categorizes existing and new NAC methods and makes them resilient to network failures by a resilient resource management. We compare the efficiency of the NAC methods with and without resilience requirements and show that they have a significant impact on the required backup capacity when resilience is required.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Real-time services such a voice over IP, IP-TV, or video conferences are getting more and more popular. They require a minimum transmission rate by the network and cannot tolerate significant packet loss or delay when congestion occurs. There are two options to satisfy their demands: providing sufficient transmission capacity [1] or limiting the traffic carried by the network at the expense of blocked flows [2]. This work focuses on the second approach which is called admission control (AC). We distinguish between link AC (LAC) methods that limit the number of flows on a link taking traffic characteristics into account and network AC (NAC) methods that limit the number of flows in a network taking the routing of the flows into account. Various NAC methods exist that differ, e.g., with respect to the number of admission control entities that need to be passed per flow [3]. An alternative classification of NAC methods is proposed in [4]. Conventional admission control avoids congestion due to increased user activity. However, in wide area networks congestion is mostly caused by traffic that has been rerouted due to link failures [5].

The contribution of this paper is twofold. We add resiliency to the concept of NAC in order to avoid congestion due to increased user activity and rerouted traffic. The configuration of non-resilient NAC implies that links cannot be overloaded by admitted traffic under failure-free operation. The configuration of resilient NAC requires that links are not overloaded by admitted traffic under failure-free operation and in likely failure scenarios. Furthermore, we compare the efficiency of different NAC methods with and without resilience requirements. The performance measure is the average resource utilization in a network that has been optimally capacitated for each NAC method based on a given routing, traffic, and target blocking probability. In the following, we clarify the relation of this paper to previous work and to other traffic engineering methods.

The performance evaluation in this study uses a capacity dimensioning approach and compares the required resources for various NAC methods under various load conditions. This is unlike real network operation where the network with capacitated links is given and NAC parameters need to be configured appropriately. This

* Corresponding author.
E-mail address: menth@informatik.uni-wuerzburg.de (M. Menth).

1389-1286/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2008.05.016

Author's personal copy

problem has been studied in [6 and 7] for networks with and without resilience requirements. Different strategies for the performance comparison of NAC methods have been discussed in [8] and the network dimensioning approach turned out to be most appropriate for that purpose. The NAC efficiency, i.e. the average resource utilization, significantly depends on the traffic characteristics. The results of [8] show how different traffic types impact the efficiency of admission control methods such that we consider only a single traffic type in this work. The efficiency of NAC methods also depends on the structure of the network and the traffic matrix. Their impact is illustrated in [9] and in [10] such that we consider only one specific network and traffic matrix in this paper. Preliminary results regarding the efficiency of resilient NAC methods have been presented in [11] while this paper gives a complete view and comparison of budget based NAC methods with and without resilience requirements. In addition, recent advances by the IETF towards a simple resilient NAC are summarized and results have been reported in [12].

Blocking due to admission control can be reduced by intelligent routing schemes. As our intention is to compare the efficiency of different NAC methods, we stick to traffic load unaware single shortest path routing and rerouting only. This assures that the differences in efficiency are due to the NAC methods and not due to other side effects. For the same reason, online routing optimization in failure cases or other advanced traffic engineering methods are out of scope in this study.

This paper is structured as follows. Section 2 gives an overview of basic network admission control (NAC) methods. Section 3 explains the performance evaluation framework which is based on capacity dimensioning. Section 4 adds resiliency to NAC and extends the framework. Section 5 compares the resource efficiency of resilient and non-resilient NAC methods. Section 6 gives a short summary and draws conclusions.

2. Methods for network admission control (NAC)

In this section, we distinguish between link and network admission control (LAC, NAC) and introduce four fundamentally different NAC concepts.

2.1. Link and network admission control

QoS criteria are usually formulated in a probabilistic way, e.g., the packet loss probability and the probability that the delay of a packet exceeds a certain threshold must both be lower than some objective values. 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 violating their QoS criteria. This includes two different aspects. Firstly, bursty traffic requires more bandwidth for transmission than its mean rate to keep the queuing delay low which can be predicted by queuing formulae [13]. The resulting capacity requirement per flow is called effective bandwidth [14]. Secondly, flows usually indicate a larger mean rate than required just to make sure that there is enough bandwidth available when needed. To take advantage of this fact, measurement based AC (MBAC) [15,16] or intentional overbooking by the provider [17] may be used. These mechanisms limit the traffic load primarily on a single link, so we call them LAC.

In contrast to LAC, network AC (NAC) coordinates the network-wide AC decisions for a flow. Thus, NAC is a distributed problem and its resource management takes the paths of the flows into account. In practice, LAC and NAC are combined: LAC calculates the effective bandwidth $c(f)$ for f which is used as input for NAC. As NAC is the focus of this work and in particular budget based NAC methods. They consist of several distributed components that possibly cooperate to admit or reject flows. We call them AC or NAC entities. Each of them has a budget b with capacity $c(b)$ and knowledge about its set of currently admitted flows $\mathscr{F}(b)$. We consider the admission process for a single AC entity. We assume that flows f are given with their effective bandwidth $c(f)$. We do not study this issue any further, and instead of pursuing hard QoS guarantees we just want to limit the overall flow rates on a link that are induced by the effective bandwidths to obtain a controlled-load service [18]. When a new flow f^{new} requests admission, it is admitted if can be accommodated by budget of the AC entity together with the already admitted flows, i.e., the following resource inequality must hold:

$$\sum_{f\in\mathscr{F}(b)} c(f) + c(f^{\text{new}}) \leqslant c(b). \tag{1}$$

In this case, the new flow is added to $\mathscr{F}(b)$, otherwise it is blocked. Upon termination, flows are removed from the sets of admitted flows. Various budget based NAC methods exist that differ with respect to the number of NAC entities in the network, the number of NAC entities by which a new flow needs to be accepted, and the type of flows that can be admitted by a specific NAC entity. In the following sections we present four major theoretic NAC methods and explain how they are implemented in practice.

2.2. Link budget based network admission control (LB NAC)

To formalize the NAC procedure, we introduce some notation. A networking scenario $\mathscr{N} = (\mathcal{V}, \mathscr{E}, u)$ is given by a set of routers $\mathcal{V}$ and set of links $\mathscr{E}$. The border-to-border (b2b) traffic aggregate with ingress router v and egress router w is denoted by $g_{v,w}$, and $\mathscr{G}$ is the set of all b2b traffic aggregates in the network. The third component of the networking scenario is the routing function u. The function $u_l(g_{v,w})$ indicates the percentage of the rate $c(g_{v,w})$ from traffic aggregate $g_{v,w}$ that is carried over link l. This notation is able to describe both single- and multi-path routing. An overview of the notation used in this paper is provided in Appendix.

The link-by-link NAC is probably the most intuitive NAC method. The capacity $c(l)$ of each link l in the network is managed by a single link budget LB_l with size $c(\text{LB}_l)$. It may be administered, e.g., at the router sending over that link or in a centralized database. A new flow $f^{\text{new}}_{v,w}$ with ingress router v, egress router w, and bit rate $c(f^{\text{new}}_{v,w})$ must pass the AC procedure for the LBs of all links that are traversed in the network by $f^{\text{new}}_{v,w}$ (cf. Fig. 1a). The NAC procedure is successful if the link budget LB_l for every link that

Author's personal copy

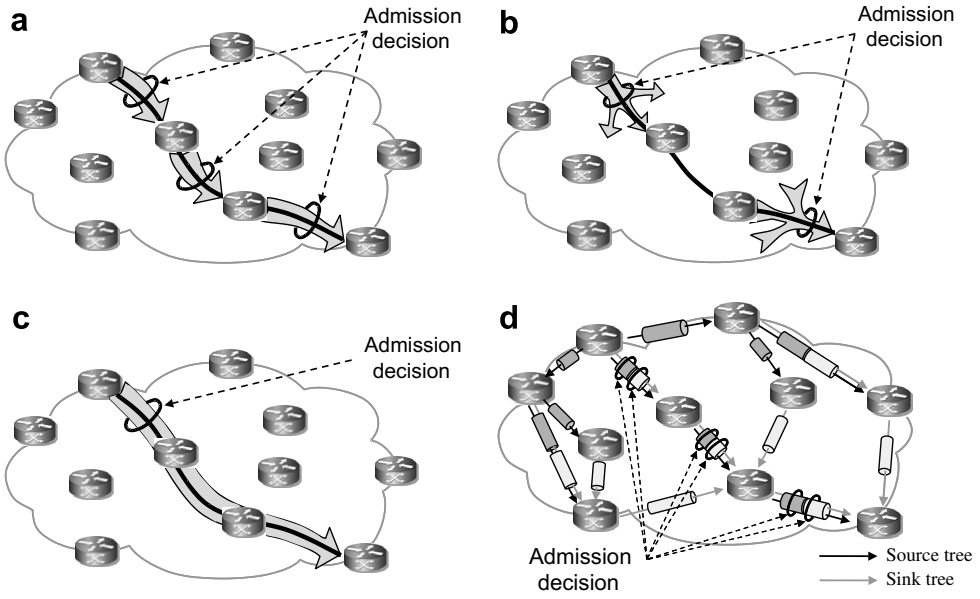

Fig. 1. The four basic network admission control (NAC) methods. (a) Link budget based NAC. There is one link budget (LB) per link. A flow must be admitted by all link budgets along its path. (b) Ingress and egress budget based NAC. There is one ingress and egress budget (IB, EB) per ingress and egress. A flow must be admitted by the appropriate ingress and egress budget. (c) Border-to-border budget based NAC. There is one border-to-border budget (BBB) per ingress–egress pair. A flow must be admitted by the appropriate border-to-border budget. (d) Ingress and egress link budget based NAC. There is one ingress and egress link budget (ILB, ELB) per link and ingress or egress node. With single-path routing the non-zero capacity budgets for an ingress or egress node form a source or sink tree. A flow must be admitted by all ingress and egress link budgets along its path.

carries traffic of the new flow can still accommodate the expected traffic share. This is expressed by the following inequality:

$$\forall l \in \mathscr{E} : u_l(g_{v,w}) > 0 : c(f_{v,w}^{\text{new}}) \cdot u_l(g_{v,w}) + \sum_{f_{x,y} \in \mathscr{F}(\text{LB}_l)} c(f_{x,y}) \cdot u_l(g_{x,y}) \leqslant c(\text{LB}_l). \quad (2)$$

with $\mathscr{F}(\text{LB}_l)$ being the set of flows that are already admitted for LB_l.[1]

There are many systems and protocols working according to that principle. The connection AC in ATM and the Integrated Services [19] architecture in IP technology adopt it in pure form and induce per flow reservation states in the core. A bandwidth broker [20–22] administers the budgets in a central database and represents a single point of failure. The stateless core approaches [23–25] work similarly with regard to their dynamic, but they avoid reservation states in the core at the expense of measurements or increased response time. In the following, we present three basic NAC methods that manage the budgets related to a flow at the flow's ingress or egress border router and keep thereby the network core truly stateless.

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. A new flow $f_{v,w}^{\text{new}}$ must pass the AC procedure for IB_v and EB_w and it is only admitted if both are successful (cf. Fig. 1b). Hence, the following inequalities must hold:

$$c(f_{v,w}^{\text{new}}) + \sum_{f \in \mathscr{F}(\text{IB}_v)} c(f) \leqslant c(\text{IB}_v), \quad (3)$$

$$c(f_{v,w}^{\text{new}}) + \sum_{f \in \mathscr{F}(\text{EB}_w)} c(f) \leqslant c(\text{EB}_w). \quad (4)$$

Thus, the NAC procedure does not require any path information about the flow. Therefore, 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 – traffic patterns that are admissible by the IBs and EBs. Hence, sufficient capacity must be provided on the links or the IBs and EBs must be set small enough to avoid congestion in the presence of certain traffic patterns.

If we leave the EBs aside, only Eq. (3) must be met for the AC procedure. This simple IB NAC originates from the DiffServ context [26,27] where traffic is admitted only at the ingress routers without looking at the destination address of the flows. The QoS of priority traffic should be guaranteed by a sufficiently low utilization of the network resources of that traffic class. To avoid any confusion: DiffServ is a mechanism for differentiated forwarding of classified traffic while IB NAC is just one NAC method that has been discussed within that context for the first time. In the same way, the other NAC approaches can also be combined with forwarding priorities for certain traffic classes, i.e. they can also be implemented in DiffServ networks.

[1] The expression $\forall x : A(x) : B(x)$ means: for all x for which the expression $A(x)$ holds, the expression $B(x)$ is true. This composed expression is either true or false.

Author's personal copy

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 procedure. A border-to-border (b2b) budget (BBB) $\mathrm{BBB}_{v,w}$ manages the capacity of a virtual tunnel between v and w and a new flow $f_{v,w}^{\mathrm{new}}$ going from ingress v to egress w passes only the AC procedure for $\mathrm{BBB}_{v,w}$ (cf. Fig. 1(c)). The AC procedure succeeds if the following inequality holds:

$$c(f_{v,w}^{\mathrm{new}}) + \sum_{f \in \mathscr{F}(\mathrm{BBB}_{v,w})} c(f) \leqslant c(\mathrm{BBB}_{v,w}). \tag{5}$$

The BBB NAC also avoids states inside the network because its budgets $\mathrm{BBB}_{v,w}$ may be controlled at the ingress or egress nodes. Every budget $\mathrm{BBB}_{v,w}$ owns a private share of the network capacity which can be used only by the specific b2b aggregate $g_{v,w}$, i.e., it cannot be used to carry other traffic with a different source or destination. Therefore, the concept is often realized in a more flexible manner, such that the size of the BBBs can be rearranged [28–30]. The tunnel capacity may be signaled using explicit reservation states in the network [31,32], only logically like in bandwidth brokers [21], or it may be assigned by a central entity [33].

2.5. Ingress link budget and egress link budget based network admission control (ILB/ELB NAC)

The ILB/ELB NAC defines ingress link budgets $\mathrm{ILB}_{l,v}$ and egress link budgets $\mathrm{ELB}_{l,w}$ to manage the capacity of each link in the network $l \in \mathscr{E}$. They are administered by border routers v and w, respectively. With single-path IP routing, traffic is forwarded along source and sink trees from ingress nodes or to egress nodes. Therefore, non-zero capacity ingress and egress link budgets of an ingress or egress node also form a logical source and sink tree when the budgets are dimensioned in a reasonable way. A new flow $f_{v,w}^{\mathrm{new}}$ must pass the AC procedure for the $\mathrm{ILB}_{l,v}$ and $\mathrm{ELB}_{l,w}$ of all links l that are traversed in the network by $f_{v,w}^{\mathrm{new}}$.

The NAC procedure will be successful if the following inequalities are fulfilled:

$$\forall l \in \mathscr{E} : u_l(g_{v,w}) > 0 : c(f_{v,w}^{\mathrm{new}}) \cdot u_l(g_{v,w}) + \sum_{f_{v,y} \in \mathscr{F}(\mathrm{ILB}_{l,v})} c(f_{v,y}) \cdot u_l(g_{v,y}) \leqslant c(\mathrm{ILB}_{l,v}) \tag{6}$$

and

$$\forall l \in \mathscr{E} : u_l(g_{v,w}) > 0 : c(f_{v,w}^{\mathrm{new}}) \cdot u_l(g_{v,w}) + \sum_{f_{x,w} \in \mathscr{F}(\mathrm{ELB}_{l,w})} c(f_{x,w}) \cdot u_l(g_{x,w}) \leqslant c(\mathrm{ELB}_{l,w}). \tag{7}$$

There are several significant differences between ILB/ELB and BBB NAC. A BBB covers only flows with the same source and destination. In contrast, ILBs cover flows with the same source but different destinations, and ELBs cover flows with the same destination but different sources. Therefore, the capacity of ILBs and ELBs can be used more flexibly than the capacity of BBBs. The BBB NAC is simpler to implement because only one $\mathrm{BBB}_{v,w}$ is checked while with ILB/ELB NAC the number of budgets to be checked is twice the flow's path length in links. Unlike LB NAC, budgets of the ILB/ELB NAC are controlled only at the border routers. Like with IB/EB NAC, there is the option to use only ILBs or ELBs by applying only Eq. (6) or Eq. (7). The concept of ILB/ELB or ILB NAC can be viewed as local bandwidth brokers at the border routers that dispose over a fraction of the network capacity. These concepts are new and have not yet been implemented by any resource management protocol. The path of the sessions in BGRP [34] matches also a sink tree, but BGRP works like LB NAC regarding its reservation dynamics.

3. Performance evaluation framework for budget based NAC methods

The objective is to compare the efficiency of various NAC methods. First we explain why we use a capacity dimensioning approach for the performance comparison. Then the traffic characteristics and the basic capacity dimensioning method used for the analysis are presented. We derive equations for the calculation of NAC-specific budget capacities and the corresponding link capacities. Finally, we define the "resource efficiency" as performance measure for the comparison in Section 5.

3.1. Approaches for the comparison of NAC METHODS

To investigate the performance of NAC methods in a network with a given topology and routing, two out of the following three parameters can be chosen and the third is determined by them such that it can serve as metric in performance evaluations: (1) the traffic matrix, (2) the network capacity, and (3) the flow blocking probability.

- Assuming (2) and (3) are given, the supportable traffic matrix can be calculated. However, the result is not well-defined and hard to compare since long flows in terms of path length require more capacity than short flows.
- Assuming (1) and (2) are given, the flow blocking probabilities for all b2b aggregate can be calculated. They yield a $|\mathcal{V}| \cdot (|\mathcal{V}| - 1)$-dimensional[2] performance measure which is not suitable for comparison purposes. A reduction to an average flow blocking probability is possible but difficult to interpret as it strongly depends on the relation between the offered traffic and the provided capacity.
- Assuming (1) and (3) are given, the required capacity $c(l)$ for every link $l \in \mathscr{E}$ in the network is calculated. The sum of all link capacities yields the network capacity $c(\mathscr{E}) = \sum_{l \in \mathscr{E}} c(l)$ which is an easy to compare and well-defined performance measure. To make the measure more intuitive, we normalize it with the average overall traffic volume in the network. This yields the average utilization of the resources that are required to achieve the desired blocking probability for all flows. Further details are given in Section 3.4.

[2] $|\mathcal{X}|$ yields the cardinality of set $\mathcal{X}$.

Author's personal copy

We use the last approach for the comparison of different NAC methods as its performance measure is intuitive and most simple to compare.

3.2. Capacity dimensioning

AC guarantees QoS for admitted flows at the expense of flow blocking if the budget capacity is exhausted. To keep the blocking probability small, the capacity $c(b)$ of a budget b must be dimensioned large enough. We review a general approach for capacity dimensioning and derive suitable budget blocking probabilities for the analysis of individual NAC methods.

3.2.1. Capacity dimensioning for a single budget

Capacity dimensioning calculates the required bandwidth for given traffic volume 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 call holding time. Although typical Internet traffic has different characteristics on the packet level [35], the Poisson model, which is used in the telephony world, is still realistic for the resource request level of user-driven real-time applications [36]. The offered load a is the mean number of active flows, provided that no flow blocking occurs. In a multi-service world like the Internet, the request profile is multi-rate, so we take n_r different request types r_i, $0 \leqslant i < n_r$ with a bitrate $c(r_i)$ and a probability of $p_r(r_i)$. 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 $E[C] = \sum_{0 \leqslant i < n_r} c(r_i) \cdot p_r(r_i) = 256$ kbit/s. The recursive solution by Kaufman and Roberts [13] allows for the computation of request type specific blocking probabilities $p_b(r_i)$ if a certain capacity c and the request type specific offered loads $a(r_i)$ are provided. We use Eq. (8) to relate the blocking probability p_b to the traffic volume instead of to the number of flows:

$$p_b = \frac{\sum_{0 \leqslant i < n_r} p_b(r_i) \cdot c(r_i) \cdot p_r(r_i)}{E[C]}. \tag{8}$$

An adaptation of the Kaufman and Roberts algorithm yields the required capacity c for a desired target blocking probability p_b [37]. Applying this principle to the traffic offered to an AC entity, 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.

3.2.2. From b2b blocking probabilities to budget blocking probabilities

Budget sizes are dimensioned using a desired budget blocking probability $p_b(b)$ which is equal to or smaller than the b2b blocking probabilities $p_{b2b}(g)$ of the flows from the traffic aggregates g using this budget b. The set $\mathscr{B}(g)$ consists of the 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

$$p_{b2b}(g) = 1 - \Pi_{b \in \mathscr{B}_g}(1 - p_b(b)) \tag{9}$$

under the assumption that flow blocking at different budgets is independent. Since flow blocking at different budgets tends to be positively correlated, this computation of $p_{b2b}(g)$ is rather conservative. In [37], we have proposed three different methods for setting the budget blocking probabilities $p_b(b)$ to achieve a desired b2b flow blocking probability $p_{b2b}(g)$. They have hardly any effect on the NAC performance, therefore, we assume a common target blocking probability $p_b(b)$ for all budgets $b \in \mathscr{B}(g)$. We denote by $m(b)$ the maximum number of budgets to be checked for any flow controlled by b. Then the required $p_b(b)$ can be derived from Eq. (9) and is $p_b(b) \leqslant 1 - \sqrt[m(b)]{1 - p_{b2b}}$.

3.3. Budget and link dimensioning for budget based NAC methods

We denote the offered load for a b2b aggregate $g_{v,w}$ by $a(g_{v,w})$ and we call $A_{\mathscr{G}} = (a(g_{v,w}))_{v,w \in \mathcal{V}}$ the traffic matrix. In contrast, the current requested rate of an aggregate is $c(g_{v,w})$ and the matrix $C_{\mathscr{G}} = (c(g_{v,w}))_{v,w \in \mathcal{V}}$ describes an instantaneous traffic pattern. A valid traffic pattern $C_{\mathscr{G}} \in \mathbb{R}_0^{+|\mathcal{V}|^2}$ obeys the following constraints:

$$\forall v, w \in \mathcal{V} : c(g_{v,w}) \geqslant 0, \tag{10}$$

$$\forall v \in \mathcal{V} : c(g_{v,v}) = 0. \tag{11}$$

If NAC is applied in the network, each traffic pattern $C_{\mathscr{G}}$ satisfies the constraints defined by the NAC budgets. Therefore, the minimum capacity $c(l)$ on link $l \in \mathscr{E}$ guaranteeing the QoS for the admitted traffic can be derived from the following link-specific worst-case analysis

$$c(l) \geqslant \max_{C_{\mathscr{G}} \in \mathbb{R}_0^{+|\mathcal{V}|^2}} \sum_{g \in \mathscr{G}} c(g) \cdot u_l(g), \tag{12}$$

when Eqs. (10) and (11) are considered together with the bandwidth constraints for the individual NAC methods. In the following, we derive these side constraints and propose efficient calculations for $c(l)$ if possible. Since the aggregate rates have real values, the maximization can be performed by the Simplex algorithm in polynomial time.

3.3.1. 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 $m(\mathrm{LB}_l) = \max_{\{g \in \mathscr{G}: u_l(g) > 0\}} len^{\max}_{\mathrm{path}}(g, l)$ whereby $len^{\max}_{\mathrm{path}}(g, l)$ is the maximum length of a (multi-)path containing l used by g. As the budget LB_l covers all flows traversing link l, its expected offered load is

$$a(\mathrm{LB}_l) = \sum_{g \in \mathscr{G}} a(g) \cdot u_l(g). \tag{13}$$

Based on $a(\mathrm{LB}_l)$, the budget capacity $c(\mathrm{LB}_l)$ is calculated using the algorithm discussed in Section 3.2.1. According to Eq. (2), the inequality

$$\forall l \in \mathscr{E} : \sum_{g \in \mathscr{G}} c(g) \cdot u_l(g) \leqslant c(\mathrm{LB}_l) \tag{14}$$

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

$$c(l) \geqslant c(\mathrm{LB}_l). \tag{15}$$

Author's personal copy

3.3.2. IB/EB NAC

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

$$a(\mathrm{IB}_v) = \sum_{w\in\mathcal{V}} a(g_{v,w}) \quad \text{and} \quad a(\mathrm{EB}_w) = \sum_{v\in\mathcal{V}} a(g_{v,w}). \tag{16}$$

We use the inequalities from Eqs. (3) and (4) as side conditions for the Simplex method in Eq. (12) to compute the capacity $c(l)$:

$$\forall v \in \mathcal{V} : \sum_{w\in\mathcal{V}} c(g_{v,w}) \leqslant c(\mathrm{IB}_v) \tag{17}$$

and

$$\forall w \in \mathcal{V} : \sum_{v\in\mathcal{V}} c(g_{v,w}) \leqslant c(\mathrm{EB}_w). \tag{18}$$

In case of the mere IB NAC, we have only $m(\mathrm{IB}_v) = 1$. The budget capacities $c(\mathrm{IB}_v)$ are computed in the same way like above, but there is a computational shortcut to calculate the required link capacity $c(l)$:

$$c(l) \geqslant \sum_{v\in\mathcal{V}} c(\mathrm{IB}_v) \cdot \sum_{w\in\mathcal{V}} u_l(g_{v,w}). \tag{19}$$

3.3.3. BBB NAC

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

$$a(\mathrm{BBB}_{v,w}) = a(g_{v,w}). \tag{20}$$

Since Eq. (5) is checked for admission,

$$\forall v, w \in \mathcal{V} : c(g_{v,w}) \leqslant c(\mathrm{BBB}_{v,w}) \tag{21}$$

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

$$c(l) \geqslant \sum_{v,w\in\mathcal{V}} c(\mathrm{BBB}_{v,w}) \cdot u_l(g_{v,w}). \tag{22}$$

3.3.4. ILB/ELB NAC

The ILB/ELB NAC requires a flow to ask ILB and ELB budgets for any link in its path for admission. Therefore, we set $m(\mathrm{ILB}_{l,v}) = 2 \cdot \max_{\{w\in\mathcal{V}:u_l(g_{v,w})>0\}}(len^{\max}_{\mathrm{path}}(g_{v,w}, l))$ and $m(\mathrm{ELB}_{l,w}) = 2 \cdot \max_{\{v\in\mathcal{V}:u_l(g_{v,w})>0\}}(len^{\max}_{\mathrm{path}}(g_{v,w}, l))$ which is similar to LB NAC. The budget $\mathrm{ILB}_{l,v}$ ($\mathrm{ELB}_{l,w}$) controls all flows with the same ingress router v (egress router w) that use link l. The offered load for these budgets is

$$a(\mathrm{ILB}_{l,v}) = \sum_{w\in\mathcal{V}} a(g_{v,w}) \cdot u_l(g_{v,w}) \tag{23}$$

and

$$a(\mathrm{ELB}_{l,w}) = \sum_{v\in\mathcal{V}} a(g_{v,w}) \cdot u_l(g_{v,w}). \tag{24}$$

Due to Eqs. (6) and (7), the side conditions

$$\forall v \in \mathcal{V} : \sum_{w\in\mathcal{V}} c(g_{v,w}) \cdot u_l(g_{v,w}) \leqslant c(\mathrm{ILB}_{l,v}) \tag{25}$$

and

$$\forall w \in \mathcal{V} : \sum_{v\in\mathcal{V}} c(g_{v,w}) \cdot u_l(g_{v,w}) \leqslant c(\mathrm{ELB}_{l,w}) \tag{26}$$

must be respected for the computation of the link capacities in Eq. (12). In case of the mere ILB NAC, the shortcut

$$c(l) \geqslant \sum_{v\in\mathcal{V}} c(\mathrm{ILB}_{l,v}) \tag{27}$$

can be applied to calculate the required link capacity if $m(\mathrm{ILB}_{l,v}) = \max_{\{w\in\mathcal{V}:u_l(g_{v,w})>0\}} len^{\max}_{\mathrm{path}}(g_{v,w}, l)$ is used to calculate the respective budget blocking probability $p_b(\mathrm{ILB}_{l,v})$.

3.4. Performance measure for NAC comparison

We compute for all NAC methods the link capacities that are required to achieve a desired b2b flow blocking probability according to the equations above. The overall required network capacity $c(\mathscr{E}) = \sum_{l\in\mathscr{E}} c(l)$ is the sum of all required link capacities in the network. The overall transmitted traffic rate $\hat{c}(\mathscr{E}) = (1 - p_{b2b}) \cdot E[C] \cdot \sum_{\{g\in\mathscr{G}\}} a(g) \cdot len^{\mathrm{avg}}_{\mathrm{path}}(g)$ is the sum of the offered load of all b2b aggregates g weighted by their average path lengths $len^{\mathrm{avg}}_{\mathrm{path}}(g)$, their acceptance probability $(1 - p_{b2b})$, and the mean rate $E[C]$ requested by a single flow. We can neglect the fact that requests with a larger rate have a higher blocking probability due to the definition of the blocking probability in Eq. (8). The overall resource utilization $\rho = \frac{\hat{c}(\mathscr{E})}{c(\mathscr{E})}$ is the fraction of the transmitted traffic rate and the overall required network capacity. We use it in Section 5 as the performance measure for the comparison of NAC methods.

4. Capacity dimensioning under resilience requirements

We present changes in the resource management that are required for resilient NAC and its performance evaluation and discuss how resilient NAC can be implemented.

4.1. Resource management for resilient NAC

Conventional NAC limits the number of flows to avoid congestion in failure-free scenarios. In case of failures, rerouted traffic may cause congestion on backup paths. In contrast, resilient NAC limits the admitted traffic to avoid congestion in failure-free scenarios and in failures cases when traffic is rerouted. That means, spare capacity must remain unallocated under failure-free condition.

Resilience can be provided only for a limited set $\mathscr{S}$ of protected failures scenarios. Each $s \in \mathscr{S}$ reflects a set of failed links and nodes. The failure-free scenario is denoted by $s^* = \emptyset$ and always contained in $\mathscr{S}$ to simplify the handling of this special case. The routing system reacts to failures by rerouting or protection switching. We describe this by the failure-specific routing function $u^s_l(g_{v,w})$.

Author's personal copy

The objective is to provide sufficient capacity $c(l)$ for each link $l \in \mathscr{E}$ such that all admissible traffic can be carried in all failure scenarios $s \in \mathscr{S}$. Hence, the minimum required link capacity $c(l)$ can be calculated by $c(l) = \max_{s\in\mathscr{S}}(c_s(l))$ where $c_s(l)$ is the required link capacity for the protected failure scenario $s \in \mathscr{S}$. We explain how $c_s(l)$ can be computed. As outlined before, NAC limits the traffic in a network by Eqs. (2)–(7). They lead to the Inequalities (14), (17), (18), (21), (25), and (26) which can be used in a linear program to evaluate the required link capacities.

As NAC entities remain unaware of the network outage in a failure case, the admission control rules in Section 2.2 remain unchanged and, therefore, the bandwidth constraints for the rate maximization in Section 3.3 also stay the same for failure scenarios, i.e., they still use the routing function $u_l^*(g_{v,w})$ for the failure-free case. However, the routing function $u_l(g_{v,w})$ is changed to $u_l^s(g_{v,w})$ in an outage scenario s. This must be respected in the traffic maximization step in Eq. (12). Due to this extension, the shortcuts for the calculation of the link capacities for the LB NAC in Eq. (15) and for the ILB NAC in Eq. (27) are not valid anymore, and the time-consuming Simplex method must be applied like for the IB/EB and for the ILB/ELB NAC. The computation shortcut for the IB NAC in Eq. (19) and for the BBB NAC in Eq. (22) can be used if $u_l(g_{v,w})$ is substituted by $u_l(g_{v,w})$.

4.2. Implementation of resilient NAC Systems

Admission decisions are usually controlled by policing entities that drop or delay packets exceeding the admitted traffic profile. The implementation of a resilient NAC requires that policing is performed only at the network border such that in case of a failure, the traffic can be redirected inside the network without being dropped by policing functions. The resilience extension of the BBB NAC, ILB NAC, ILB/ELB NAC, IB NAC, and IB/EB NAC are straightforward because their budgets are controlled at the network border and just need to be set low enough that the admitted traffic can be carried inside the network under normal conditions and in protected failure scenarios $\mathscr{S}$.

This is different with LB NAC. An implementation of a resilient LB NAC is currently discussed in the IETF in the context of pre-congestion notification (PCN) [38]. The standardization is still in an early phase and, therefore, the nomenclature is not stable yet. Each link $l \in \mathscr{E}$ in the PCN domain is associated with an admissible rate $AR(l)$. If the PCN traffic exceeds this threshold, all packets on that link are re-marked with an "admission-stop" codepoint. The egress nodes monitor the markings of the packets and notify the ingress nodes from which they have received the marked packets to stop admission of new flows. As a consequence, flows can be controlled by AC and policing entities at the network border, although the capacity of the link budgets inside the network is not under the control of a single AC entity. Resilient LB NAC can be implemented by this mechanism by setting the admissible rate of a link to $AR(l) = c(\mathrm{LB}_l)$.

The PCN framework also comprises a flow termination function in the presence of severe congestion which can be caused by rerouting in unprotected failure scenarios or by malfunction of the AC due to its measurement based nature. This flow termination function also allows to use relaxed budget capacities that are resilient for likely traffic patterns. When unlikely traffic patterns coincide with a failure and congestion occurs, the overload can be resolved by terminating traffic. This leads to a more efficient PCN based NAC [12].

5. Performance of NAC methods with and without resilience requirements

In this section, we present numerical results for the resource utilization of all NAC methods with and without resilience requirements and discuss them in detail. The results are obtained analytically based on the equations in the previous sections. In all our studies we dimension the network capacity to meet a flow blocking probability of $p_{b2b} = 10^{-3}$. We have shown that the offered load has the major impact on the required capacity and the resource utilization [37]. Therefore, we consider the performance only depending on the offered load a_{b2b}, which is the mean load between two border routers. In our investigation, we use the COST-239 core network (cf. Fig. 2, [39]), a homogeneous traffic matrix, and single shortest path routing as it is used by IS–IS or OSPF. For the sake of completeness, the impact of different blocking probabilities, other network topologies, heterogeneous traffic matrices, and different routing is studied in [37].

5.1. Performance of NAC methods without resilience requirements

If the offered load is large, the capacity required to meet a given flow blocking probability can be better utilized than if the offered load is low. This observation is called economy of scale or multiplexing gain, and it is the key for understanding NAC performance. Fig. 3 shows the performance of all NAC types for single-path (SP) routing without resilience requirements. We observe the typical increase of the resource utilization with the offered b2b load a_{b2b}. The LB, ILB/ELB, ILB, and BBB NAC can achieve 100% resource utilization in the limit. The IB/EB NAC has a better performance than the IB NAC, but they are both inefficient as their curves converge to network topology specific asymptotes of 22% and 16%, respectively.

The differences among the efficient NAC types result from their different ability to realize multiplexing gain. The link budgets cover the largest amount of traffic (cf. Eq. (13)), followed by ingress and egress link budgets (cf. Eqs. (23) and (24)), and by b2b budgets (cf. Eq. (20)). The increased offered load leads to more multiplexing gain and explains the order of efficiency for the LB, ILB, and BBB NAC.

The IB NAC is not economical. The budget capacity $c(\mathrm{IB}_v)$ can be used for any flow entering the network at border router v which is on the one hand very flexible, but on the other hand, it allows also very pathological traffic patterns to be accepted. E.g., an unlikely traffic pattern with $c(g_{v,w}) = c(\mathrm{IB}_v)$ can be accepted. Thus, the full budget capacity $c(\mathrm{IB}_v)$ must be provided on the links on the paths from v to any other destination w. This capacity can be

Author's personal copy

used only by the traffic that is admitted by the budget IB_v. In case of single-path routing, the source tree formed by the paths originating at any node v comprises exactly $|\mathcal{V}|-1$ links. As a consequence, the average resource utilization of the IB NAC cannot exceed a value of $\frac{len_{path}^{avg}}{|\mathcal{V}|-1} = \frac{1.56}{10} = 15.6\%$ for a homogeneous traffic matrices. The application of additional egress budgets EB_w excludes some unlikely traffic patterns from being accepted, e.g., the one where the full rate of $c(IB_v)$ streams from all ingress routers v to the same egress router w. Therefore, the IB/EB NAC limits the inefficiency to a certain extent, but it does not solve the basic problem. The ILB/ELB NAC also improves the performance of the ILB NAC by applying additional egress link budgets.

5.2. Performance of NAC methods with resilience requirements

A network $\mathcal{N} = (\mathcal{V}, \mathcal{E})$ can face $\binom{|\mathcal{E}|}{k}$ different failure scenarios with k link failures. However, the probability of an outage decreases with the number of the simultaneously failed components. Single link failures are most important [5], therefore, we restrict the set of protected failure scenarios $\mathcal{S}$ to all single bi-directional link failures. The routing in a failure scenario s adapts to the new topology according to the shortest path algorithm and provides the failure-specific routing function $u_l^s(g_{v,w})$.

Fig. 4 shows the resource utilization for all NAC methods under resilience requirements depending on the average offered b2b load. It reveals a completely different performance behavior compared to the resource utilization without resilience requirements (cf. Fig. 3). All NAC types have different asymptotes for their resource utilization and these asymptotes are network- and routing-specific [11,37]. The BBB NAC outperforms the ILB/ELB NAC, the ILB NAC, and the LB NAC. Except for ILB NAC and ILB/ELB NAC, this is the reversed order of the scenario without resilience. The performance of the IB and IB/EB NAC is again significantly worse.

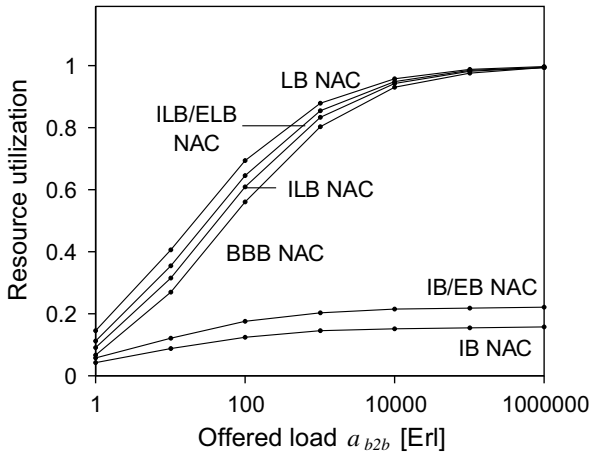

Fig. 3. Resource utilization without resilience requirements.

With resilience requirements, the BBB NAC achieves only 60% resource utilization in the limit instead of 100% without resilience requirements. The reciprocal value $\frac{1}{0.6} \approx 1.67$ is the average degree of capacity overprovisioning. It is required to maintain QoS in all protected failure scenarios and corresponds to 67% additional backup capacity. This shows that rerouting on the network layer requires less backup capacity compared to 100% on the physical layer. This rather low value is achieved since the backup capacity on the network layer can be shared by different traffic aggregates in different failure scenarios. In addition, the backup capacity on the network layer may be used to carry low priority traffic during failure-free operation which makes the restoration option even more attractive [40].

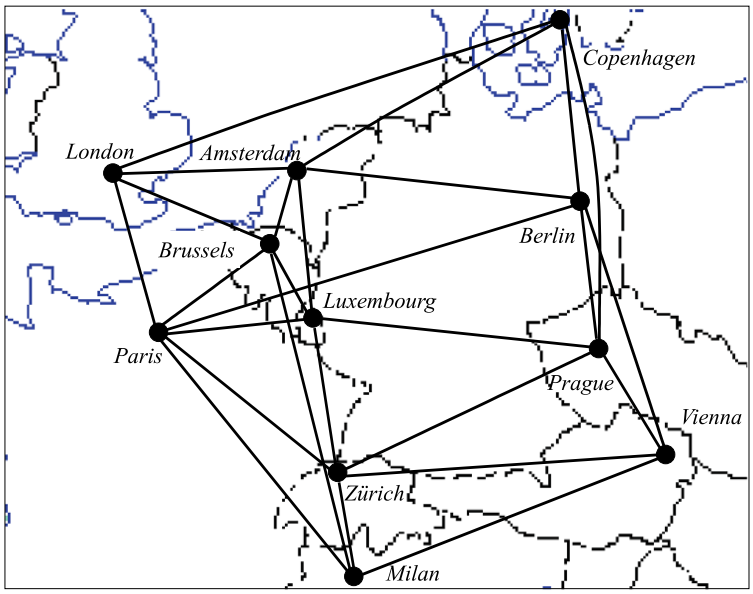

Fig. 2. The topology of the COST-239 core network.

Author's personal copy

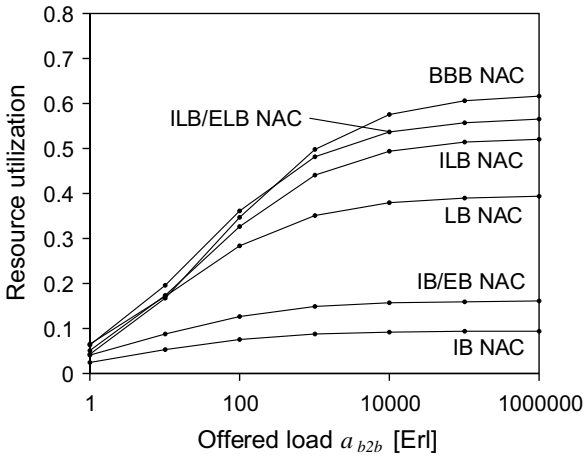

Fig. 4. Resource utilization for SP routing with resilience requirements.

The amount of required backup capacity depends on the routing in the failure-free case and in failure cases and can be minimized by optimizing the routing system [41,42]. However, when AC is applied, the required backup capacity also depends on the NAC method. With resilience requirements, the maximum resource utilization for the LB NAC is 40% which corresponds to 150% additional costs for backup purposes. Hence, the LB NAC is clearly more expensive than the BBB, ILB/ELB, and ILB NAC from a resource point of view. In addition, it is not able to offer cheap resilience for QoS services below 100% backup capacity although the routing is the same as in the experiment with the BBB NAC.

There is an explanation for that phenomenon. The LB NAC is more flexible than the BBB NAC with regard to the use of allocated link capacities, i.e., more traffic patterns can be supported with the same capacity. On the one hand, less capacity suffices to obtain the same QoS level and the LB NAC has a better resource efficiency than the BBB NAC in the non-resilient case. On the other hand, this flexibility is a drawback with resilience requirements since all admissible traffic patterns must be protected. Only little information is known about the traffic that will be admitted by an LB NAC entity, but the backup resources are allocated in the network a priori. As a consequence, backup resources need to be reserved for all admissible traffic patterns. Therefore, the reserved backup capacity for a specific failure scenario cannot be used entirely by the admitted flows at the same time. This is different with the BBB NAC since the $BBB_{v,w}$ controls only the traffic of a single traffic aggregate $g_{v,w}$ whose backup paths is known in advance for any protected failure scenario. Hence, the backup capacity can be reserved in advance in a more focussed way and it can be fully used in the failure case by the admitted traffic. In a nutshell, a large NAC flexibility with regard to traffic patterns achieves a high resource utilization without resilience requirements, but it requires much additional capacity for backup purposes and causes a low resource utilization with resilience requirements.

In our experiment, the resource efficiency of the ILB NAC is 52% in the limit which corresponds to 92% backup capacity. It is worse than the one of the BBB NAC since the ILB NAC is more flexible than the BBB NAC, but it is better than the one of the LB NAC since the ILB NAC is less flexible than the LB NAC. The ILB/ELB NAC applies additional egress link budgets which leads to a sharper profile of the admissible traffic compared to the one of the ILB NAC. This improves the resource utilization of the ILB/ELB NAC to 56% which corresponds to 79% backup capacity.

We also observe a decrease of the resource utilization of the IB/EB NAC from 22% without resilience requirements to 16% with resilience requirements. The additional expenses for backup purposes are only 37.5%, but the absolute required network capacity exceeds the demand of the other NAC methods by far and leaves the IB/EB still unattractive. The same holds for the IB NAC with a maximum resource utilization of 9% opposed to 16% without resilience requirements.

6. Conclusion

In this paper, we proposed the concept of resilient network admission control (NAC). We extended the resource management for four fundamentally different NAC types in such a way that admitted traffic cannot cause congestion when it is rerouted due to a protected failure. We compared the efficiency of the NAC methods with and without resilience requirements. The performance measure is the average resource utilization in an optimally dimensioned network.

The direct comparison of the NAC methods without resilience requirements showed that the LB NAC is most efficient, followed by the ILB/ELB NAC, the ILB NAC, and the BBB NAC. However, all these NAC types achieve a resource utilization close to 100% for sufficiently high offered load. In contrast, the average resource utilization of the IB NAC and the IB/EB NAC converges to network-specific asymptotes of 16% and 22% in the COST-239 network. Under resilience requirements, the efficient NAC methods achieve a maximum resource utilization between 40% and 67%. They have different utilization limits and the order of their efficiency is reversed, i.e., the BBB NAC is most efficient and the LB NAC is least efficient. Hence, the NAC methods have a tremendous impact on the required backup capacity. We observed the same effects in different network topologies [37] which underlines the general nature of our findings.

If NAC is deployed in IP networks, it should be resilient against failures which is technically feasible and necessary since most congestion results from rerouted traffic. Currently, the IETF is about to standardize resilient NAC based on pre-congestion notification (PCN). It is especially attractive because of its simplicity and opens a wide field for new control structures.

Acknowledgements

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.

Author's personal copy

Appendix

See Table 1.

Table 1
Overview of applied notation

$v, w \in \mathcal{V}$	Set of nodes in a network
$l \in \mathcal{E}$	Set of links in a network
$s \in \mathcal{S}$	Set of protected failure scenarios, s is a set of failed elements
b	AC budget, can be of type *LB*, *BBB*, *IB*, *EB*, *ILB*, or *ELB*
f ($f_{v,w}$)	Individual flow (from ingress node v to egress node w)
$\mathcal{F}$ ($\mathcal{F}(b)$)	Set of flows (admitted by budget b)
g ($g_{v,w}$)	Traffic aggregate (from ingress node v to egress node w)
$\mathcal{G} = \{g_{v,w} : v, w \in \mathcal{V}\}$	Set of all traffic aggregates in the network
$\mathcal{B}(g)$	Set of budgets that need to be passed to admit flows of aggregate g
$c(l)$	Bandwidth of link l
$c(b)$	Capacity of NAC budget b
$c(f)$	Effective bandwidth of flow f
$c(g)$	Current effective bandwidth of traffic aggregate g
$C_{\mathcal{G}} = (c(g_{v,w}))_{v,w \in \mathcal{V}}$	Current traffic pattern
$u_l(g_{v,w})$	Percentage of traffic from v to w that is carried over l
$u_l^s(g_{v,w})$	Routing function in the presence of failure scenario s
r_i	Request type
$p_r(r_i)$	Probability of request type r_i
$c(r_i)$	Rate of request type r_i
$a(b)$	Offered load for a budget b in Erlang
$p_b(b)$	Blocking probability at a budget b
$p_{b2b}(g)$	b2b blocking probability for flows of aggregate g
$m(b)$	Maximum number of budgets that need to be checked for a flow Which is controlled by budget b
$a(g)$	Offered load for aggregate g in Erlang
$A_{\mathcal{G}} = (a(g_{v,w}))_{v,w \in \mathcal{V}}$	Traffic matrix
$len_{\text{path}}^{\max}(g, l)$	Maximum length of a (multi-)path used by g and containing l
$\lvert\mathcal{X}\rvert$	Cardinality of set $\mathcal{X}$

References

[1] M. Menth, R. Martin, J. Charzinski, Capacity overprovisioning for networks with resilience requirements, in: ACM SIGCOMM, Pisa, Italy, 2006.
[2] S. Shenker, Fundamental design issues for the future internet, IEEE Journal on Selected Areas in Communications 13 (7) (1995) 1176–1188.
[3] M. Menth, S. Kopf, J. Milbrandt, J. Charzinski, Introduction to budget-based network admission control methods, in: 28th IEEE Conference on Local Computer Networks (LCN), Bonn, Germany, 2003.
[4] A. Riedl, T. Bauschert, J. Frings, On the dimensioning of voice over IP networks for various call admission control schemes, in: 18th International Teletraffic Congress (ITC), Berlin, Germany, 2003, pp. 1311–1320.
[5] S. Iyer, S. Bhattacharyya, N. Taft, C. Diot, An Approach to Alleviate Link Overload as Observed on an IP Backbone, IEEE Infocom, San Francisco, CA, 2003.
[6] M. Menth, S. Gehrsitz, J. Milbrandt, Fair assignment of efficient network admission control budgets, in: 18th International Teletraffic Congress (ITC), Berlin, Germany, 2003, pp. 1121–1130.
[7] M. Menth, J. Milbrandt, S. Kopf, Capacity assignment for NAC budgets in resilient networks, in: International Telecommunication Network Strategy and Planning Symposium (Networks), Vienna, Austria, 2004, pp. 193–198.
[8] M. Menth, S. Kopf, J. Milbrandt, A performance evaluation framework for network admission control methods, in: IEEE Network Operations and Management Symposium (NOMS), Seoul, South Korea, 2004, pp. 307–320.
[9] M. Menth, S. Kopf, J. Charzinski, Impact of network topology on the performance of network admission control methods, in: IEEE International Workshop on Multimedia Interactive Protocols and Systems (MIPS), Naples, Italy, 2003, pp. 195–206.
[10] M. Menth, J. Milbrandt, S. Kopf, Impact of routing and traffic distribution on the performance of network admission control, in: Ninth IEEE Symposium on Computers and Communications (ISCC), Alexandria, Egypt, 2004, pp. 883–890.
[11] M. Menth, S. Kopf, J. Charzinski, Network admission control for fault-tolerant QoS provisioning, in: IEEE High-Speed Networks for Multimedia Communication (HSNMC), Toulouse, France, 2004, pp. 1–13.
[12] M. Menth, Efficiency of PCN-based network admission control with flow termination, Praxis der Informationsverarbeitung und Kommunikation (PIK) 30 (2) (2007) 82–87.
[13] J. Roberts, U. Mocci, J. Virtamo, Broadband Network Teletraffic – Final Report of Action COST 242, Springer, Berlin, Heidelberg, 1996.
[14] F.P. Kelly, Stochastic Networks: Theory and Applications, vol. 4, Oxford University Press, 1996 (Ch. Notes on Effective Bandwidths, pp. 141–168).
[15] L. Breslau, E.W. Knightly, S. Shenker, H. Zhang, Endpoint admission control: architectural issues and performance, in: ACM SIGCOMM, 2000.
[16] L. Breslau, S. Jamin, S. Shenker, Comments on the Performance of Measurement-Based Admission Control Algorithms, IEEE Infocom, 2000, pp. 1233–1242.
[17] M. Menth, J. Milbrandt, S. Oechsner, Experience-based admission control (EBAC), in: Ninth IEEE Symposium on Computers and Communications (ISCC), Alexandria, Egypt, 2004, pp. 903–910.
[18] J. Wroclawski, RFC2211: Specification of the Controlled-Load Network Element Service, September, 1997.
[19] B. Braden, D. Clark, S. Shenker, RFC1633: Integrated Services in the Internet Architecture: An Overview, June, 1994.
[20] A. Terzis, J. Wang, J. Ogawa, L. Zhang, A two-tier resource management model for the Internet, in: Global Internet Symposium at IEEE Globecom, 1999.
[21] B. Teitelbaum, S. Hares, L. Dunn, V. Narayan, R. Neilson, F. Reichmeyer, Internet2 QBone: Building a Testbed for Differentiated Services, IEEE Network Magazine.
[22] Z.-L. Zhang, Z. Duan, Y.T. Hou, On scalable design of bandwidth brokers, IEICE Transaction on Communications E84-B (8) (2001) 2011–2025.
[23] I. Stoica, H. Zhang, Providing guaranteed services without per flow management, in: ACM SIGCOMM, Boston, MA, 1999.
[24] S. Bhatnagar, B. Nath, Distributed Admission Control to Support Guaranteed Services in Core-Stateless Networks, IEEE Infocom, San Francisco, USA, 2003.
[25] R. Szábó, T. Henk, V. Rexhepi, G. Karagiannis, Resource management in differentiated services (RMD) IP networks, in: International Conference on Emerging Telecommunications Technologies and Applications (ICETA 2001), Kosice, Slovak Republic, 2001.
[26] S. Blake, D.L. Black, M.A. Carlson, E. Davies, Z. Wang, W. Weiss, RFC2475: An Architecture for Differentiated Services, December, 1998.
[27] X. Xiao, L.M. Ni, Internet QoS: A big picture, IEEE Network Magazine 13 (2) (1999) 8–18.
[28] T. Engel, E. Nikolouzou, F. Ricciato, P. Sampatakos, Analysis of adaptive resource distribution algorithms in the framework of a dynamic DiffServ IP network, in: Eighth International Conference on Advances in Communications and Control, Crete, Greece, 2001.
[29] H. Fu, E. Knightly, Aggregation and scalable QoS: A performance study, in: IEEE International Workshop on Quality of Service (IWQoS), Karlsruhe, Germany, 2001.
[30] J. Milbrandt, M. Menth, S. Kopf, Adaptive bandwidth allocation for wide area networks, in: 19th International Teletraffic Congress (ITC), Beijing, China, 2005, pp. 1235–1244.
[31] F. Baker, C. Iturralde, F. Le Faucheur, B. Davie, RFC3175: Aggregation of RSVP for IPv4 and IPv6 Reservations, September, 2001.
[32] D.O. Awduche, L. Berger, D.-H. Gan, T. Li, V. Srinivasan, G. Swallow, RFC3209: RSVP-TE: Extensions to RSVP for LSP Tunnels, December, 2001.
[33] P. Trimintzios, I. Andrikopoulos, G. Pavlou, P. Flegkas, D. Griffin, P. Georgatsos, D. Goderis, Y. T'Joens, L. Georgiadis, C. Jacquenet, R. Egan, A management and control architecture for providing IP differentiated services in MPLS-based networks, IEEE Communications Magazine 39 (5) (2001) 80–88.
[34] P. Pan, H. Schulzrinne, BGRP: a tree-based aggregation protocol for inter-domain reservations, Journal of Communications and Networks 2 (2) (2000) 157–167.

Author's personal copy

[35] V. Paxson, S. Floyd, Wide-area traffic: the failure of poisson modeling, IEEE/ACM Transactions on Networking 3 (3) (1995) 226–244.
[36] J.W. Roberts, Traffic theory and the Internet, IEEE Communications Magazine 1 (3) (2001) 94–99.
[37] M. Menth, Efficient admission control and routing in resilient communication networks, Ph.D. thesis, University of Würzburg, Faculty of Computer Science, Am Hubland, July, 2004.
[38] P. Eardley (ed.), Pre-congestion notification architecture, February, 2008, <http://www.ietf.org/internet-drafts/draft-ietf-pcn-architecture-03.txt>.
[39] P. Batchelor et al., Ultra-High Capacity Optical Transmission Networks. Final Report of Action COST 239, January, 1999.
[40] C. Hoogendoorn, K. Schrodi, M. Huber, C. Winkler, J. Charzinski, Towards carrier-grade next generation networks, in: International Conference on Communication Technology (ICCT), Beijing, China, 2003.
[41] B. Fortz, M. Thorup, Robust optimization of OSPF/IS-IS weights, in: International Network Optimization Conference (INOC), Paris, France, 2003, pp. 225–230.
[42] A. Sridharan, R. Guerin, Making IGP routing robust to link failures, in: IFIP-TC6 Networking Conference (Networking), Ontario, Canada, 2005.

Michael Menth studied computer science and mathematics at the Universitiy of Wuerzburg/Germany and Austin/Texas. He worked at the University of Ulm/Germany and Wuerzburg and obtained his Ph.D., in 2004. Currently, he is assistant professor and heading the research group "Next Generation Networks" at the Institute of Computer Science in Wuerzburg. His special interests are performance analysis, optimization of communication networks, resource management, resilience issues, and Future Internet. Dr. Menth holds numerous patent applications and received various scientific awards for innovative work.

Stefan Kopf studied computer science at the University of Wuerzburg/Germany. He has been part of the next generation networks group of Dr. Menth during his diploma thesis and holds 2 patent applications. Currently, he is the CEO of a privately held company and works in the implementation and consulting of enterprise content management systems for large scaled enterprises.

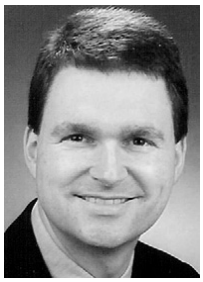

Joachim Charzinski received his Dipl.-Ing. and Dr.-Ing. degrees from University of Stuttgart, Germany, in 1991 and 1999, respectively. In 1997 he joined the public communications group of Siemens AG, Munich, Germany, where he held several positions in research and innovation management. He is currently a principal innovator at Nokia Siemens Networks. His focus topics are traffic engineering, traffic measurement and modelling, network resilience and network security. He is author or co-author of 45 journal and conference papers and holds 24 patents.

Karl J. Schrodi received his Dipl.-Ing. degree in Electrical Engineering from the University of Stuttgart, Germany, and joined the Standard Elektrik Lorenz AG (SEL) Research Center in Stuttgart, Germany, in 1980. Following the acquisition of SEL by Alcatel in 1987 he was part of Alcatel's central broadband and ATM program and headed the System Engineering department of Alcatel SEL's broadband design center. In 1996, he joined the communications group of Siemens AG which recently has been merged into Nokia Siemens Networks. Karl authored and co-authored many publications and is a two-times honoree of Inventor of the Year awards in Siemens.