

Capacity Overprovisioning for Networks with Resilience Requirements

Michael Menth
menth@informatik.uni-
wuerzburg.de

Rüdiger Martin
martin@informatik.uni-
wuerzburg.de

Joachim Charzinski
joachim.charzinski@
siemens.com

Department of Distributed Systems
Institute of Computer Science
University of Würzburg, Germany

Siemens AG
Munich, Germany

ABSTRACT

This work focuses on capacity overprovisioning (CO) as an alternative to admission control (AC) to implement quality of service (QoS) in packet-switched communication networks. CO prevents potential overload while AC protects the QoS of the traffic during overload situations. Overload may be caused, e.g., by fluctuations of the traffic rate on a link due to its normal stochastic behavior (a), by traffic shifts within the network due to popular contents (b), or by redirected traffic due to network failures (c). Capacity dimensioning methods for CO need to take into account all potential sources of overload while AC can block excess traffic caused by (a) and (b) if the capacity does not suffice. The contributions of this paper are (1) the presentation of a capacity dimensioning method for networks with resilience requirements and changing traffic matrices, (2) the investigation of the impact of the mentioned sources of overload (a-c) on the required capacity for CO in networks with and without resilience requirements, and (3) a comparison of this required capacity with the one for AC. Our results show that in the presence of strong traffic shifts CO requires more capacity than AC. However, if resilience against network failures is required, both CO and AC need additional backup capacity for the redirected traffic. In this case, CO can use the backup capacity to absorb other types of overload. As a consequence, CO and AC have similar bandwidth requirements. These findings are robust against the network size.

Categories and Subject Descriptors

C.4 [Computer-Communication Networks]: Performance of Systems

General Terms

Design, Economics, Performance, Reliability

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'06, September 11–15, 2006, Pisa, Italy.
Copyright 2006 ACM 1-59593-308-5/06/0009 ...\$5.00.

Keywords

QoS, admission control, capacity overprovisioning

1. INTRODUCTION

Quality of service (QoS) in packet-switched communication networks can be expressed in terms of limited packet loss and delay. It can be achieved by avoiding overload in networks which occurs if the traffic rate for a link exceeds a certain threshold. There are two main approaches for this objective: capacity overprovisioning (CO) and admission control (AC) [33]. CO provides so much capacity on the links that overload is unlikely while AC limits the number of flows over a specific network element, e.g. a link, to avoid overload. CO is applied in today's core networks and it is currently favored by many Internet service providers (ISPs) and researchers as the preferred mechanism for QoS [28]. However, from a scientific point of view, it is less understood than AC.

ISPs are usually in favor of CO since it keeps the networks simple while AC is complex and requires a significant amount of interoperability. In contrast, telecommunication providers usually prefer AC since they must care for guaranteed QoS which can be difficult and costly with CO. As a consequence, the discussion between both parties regarding the question which approach should be taken in an QoS-enabled Internet resembles an almost religious war [4, 14]. In this work, we contribute to this discussion by quantifying and comparing the required capacity for CO and AC under potential overload and resilience requirements.

Overload in networks may be caused, e.g., by fluctuations of the traffic rate on a link due to its normal stochastic behavior (a), by traffic shifts within the network due to popular contents (b), or by redirected traffic due to network failures (c). Capacity dimensioning methods for CO need to take into account all potential sources of overload. In contrast, AC can block excess traffic if overload occurs that is caused by (a) and (b). However, classical AC cannot guarantee QoS in the presence of overload caused by redirected traffic which is the most frequent reason for overload in core networks [19]. As a consequence, classical AC is not appropriate to guarantee QoS during failure scenarios. Resilient AC heals this shortcoming by admitting traffic only if it can be carried without QoS violation together with the redirected traffic of potential failure scenarios [21]. Thus, compared to classical AC, resilient AC needs additional backup capacity to accommodate the same traffic.

The contributions of this paper are (1) the presentation of a capacity dimensioning method for networks with resilience requirements and changing traffic matrices, (2) the investigation of the impact of the mentioned sources of overload (a-c) on the required capacity for CO in networks with and without resilience requirements, and (3) a comparison of this required capacity with the one for AC. The parametrization of the presented overload models based on traffic measurements is not topic of this paper.

This work is structured as follows. Section 2 summarizes related work regarding CO and AC. Section 3 reviews dimensioning methods for CO and AC and illustrates their impact. Section 4 presents models for simple and complex traffic shifts. It introduces and compares capacity dimensioning for CO and AC in networks with and without resilience requirements. Finally, Section 5 concludes this work.

2. RELATED WORK

In this section, we start with fundamentals of resilient networks and review then existing literature on CO, AC, and the comparison of both.

2.1 Network Resilience

Networks are protected against failures by protection and restoration mechanisms [1]. Restoration is the setup of backup paths after a failure occurred, e.g., by the reconvergence of shortest path routing. In contrast, protection sets up backup paths in advance to redirect the traffic in failure cases which is done, e.g., by MPLS fast reroute [25]. The QoS of the traffic may be violated during the failover time which is shorter for protection mechanisms than for restoration mechanisms. After that time, a smooth operation of the network is desired which demands sufficient capacity on the backup paths. In other words, redirected traffic may lead to overload when network links are not properly dimensioned or if the resilience mechanisms are not properly configured.

2.2 Capacity (Over-)Provisioning

Capacity overprovisioning (CO) relies on providing sufficient bandwidth to make overload in networks unlikely and to achieve thereby the desired QoS. The link capacities are chosen in such a way that they are very rarely exceeded by the predicted traffic.

The rate of Internet flows is often hard to determine due to long range dependency which has been found first in local area networks [18] but then also for general WWW traffic [9] and for variable bit rate realtime traffic [2]. The measurement results in [42] show that traffic fluctuations at small time scales tend to be rather uncorrelated while they reveal a self-similar structure on a time scale of 1 s and above. This is problematic since self-similarity provokes significant packet loss with finite buffers [12]. However, a superposition of sufficiently many long range dependent traffic sources may lead for moderate utilization values in the limit to the same buffer overflow probability as a Poisson process [7].

Bandwidth provisioning procedures differ fundamentally from access to core networks due to the degree of traffic aggregation. In [11] empirical evidence can be found that core network traffic on the packet level, i.e. the average traffic arrival rate, is modelled well by the Gaussian distribution due to the high level of aggregation. This is clearly not the case in the access due to the limited number of users where the aggregation level is inherently low. The network in [11]

is dimensioned to support latency sensitive traffic. Accordingly, the QoS measure the network is dimensioned for is the probability that the queue length Q of a router exceeds a certain value x : $\Pr(Q > x)$. To satisfy end-to-end delay requirements as low as 3 ms requires only 15% extra bandwidth above the average data rate of the traffic in the highly aggregated Sprint network.

Rate measurements by SNMP are usually obtained only on a time scale of 5 min and reveal substantially smaller variations than traffic on a small time scale like 10 ms. The difference of this variation may be 100% or more [38]. This makes the prediction of the required capacity based on traffic measurements very hard. The work in [39] focuses on the probability that the amount of traffic $A(T)$ generated on a link within a specified time interval T exceeds the capacity C of the link: $\Pr(A(T) \geq C \cdot T)$. The authors argue that applications can cope with lack of bandwidth within an application-dependent small interval T if this occurs sufficiently rarely. They develop an interpolation formula that predicts the bandwidth requirement on a relatively short time scale in the order of 1 s by relying on coarse traffic measurements. So-called ‘user-oriented’ and ‘black box’ traffic models are used to characterize measurement results. They are evaluated in [37] with regard to their accuracy for link provisioning. It turned out that black box models are easier to estimate and yield accurate provisioning guidelines.

Another closely related problem is forecasting of Internet traffic. A recent approach for long-term forecasting can be found in [26]. The authors of [36] combine both tasks to yield an adaptive bandwidth provisioning algorithm. Based on measurements, the required capacity is predicted and adjusted on relatively small time scales between 4 s and 2 min. The Maximum Variance Asymptotic (MVA) [8] approach for the tail probability of a buffer fed by an input Gaussian process is used to make the QoS requirement $\Pr(\text{delay} > D) < \epsilon$ explicit.

2.3 Admission Control

Admission control (AC) limits the number of flows in the network by denying access to new ones if the network risks to be overloaded. It has been proposed for the Internet in [33]. Admission control mechanisms pursue two objectives. On the one hand, the challenge on a single link is to decide whether the admission of new flow compromises the QoS in terms of packet loss and delay on that link. On the other hand, the challenge in a network is to decide whether the admission of a new flow violates the QoS on any link of its path. Many different methods and protocols have been proposed to solve both aspects which we call link AC (LAC) and network AC (NAC) in the following [21]. Implementations always solve both issues, even if one of them is implemented in a trivial way.

2.3.1 Link Admission Control

Link AC (LAC) concentrates on a single resource and primarily on the packet level. The methods can be roughly subdivided into descriptor based, measurement based, and hybrid LAC mechanisms.

2.3.1.1 Descriptor Based LAC and Effective Bandwidths.

With descriptor based LAC methods, connection requests carry a flow descriptor that is taken into account for the

admission decision. The descriptor typically characterizes the rate and the variability of a flow on different time scales. This may be done by a single or dual token bucket which includes mean and peak rates. Policers and spacers may be used to enforce the conformance of the flow characterization on input and output interfaces of a router. Mechanisms differ regarding the flow description, the calculation for the AC decision, and additional assumptions that are taken for the calculation [29]. A generalization computes the so-called effective bandwidth [16] of a flow based on its descriptor, the link bandwidth, and the QoS requirements. This is then used as an additive value to decide whether free bandwidth for a new flow is still available on the link.

2.3.1.2 *Measurement Based LAC and Hybrid Methods.*

Measurement based AC uses measurements to determine either the bandwidth requirements of individual admitted flows [15] or of the admitted traffic aggregate [13, 17] and derives thereby the free bandwidth for additional flows. The hybrid approach in [22] works like descriptor based LAC, but determines a feasible degree of overbooking which is obtained through measurement experience from the past.

2.3.2 *Network Admission Control*

Network AC (NAC) prevents overload on multiple resources within a network. This is a non-trivial task if the decision should be made solely at the network border without cooperation of interior nodes since flows take different paths that use different links of the network.

2.3.2.1 *Link-by-Link NAC and Bandwidth Brokers.*

The simplest NAC implementation is certainly the application of LAC for each link along the path of a flow. The reservation for the flow is only admitted if all AC decisions succeed. This requires interior nodes of a network to keep per flow states which is difficult to handle, in particular when network failures occur. As an alternative, the distributed actions can be performed by a central bandwidth broker that has a complete view of all network resources [24, 41]. However, the bandwidth broker constitutes a single point of failure.

2.3.2.2 *Feedback Based NAC.*

Several protocols [3, 34, 35] work according to a principle that we call feedback based NAC. The traffic sources send capacity tickets in regular intervals along the paths of the admitted flows. Interior nodes collect them and account for the reserved capacity within the last interval without knowing the individual flows. A new flow injects a capacity ticket into the network which is discarded by an interior router if overload may occur through the admission of this flow. If the capacity ticket arrives at the destination, it is returned to the issuing router which signals the admission of the reservation request; otherwise it is rejected by a timeout.

2.3.2.3 *B2B Budget Based NAC.*

The border-to-border (b2b) budget (BBB) based NAC defines capacity budgets for each b2b relationship (v, w) within the network and assigns them a capacity portion. A new flow at ingress router v and destined for egress router w requests for admission only at its ingress router v . This ingress router performs AC based on the a priori dedicated capacity

budget $BBB(v, w)$ like on a single resource. This concept is implemented, e.g., by label switched paths (LSPs) in MPLS.

2.3.2.4 *Resilient NAC.*

As mentioned above, overload occurs mainly due to network failures and redirected traffic [19], but classical NAC approaches cannot cope with these situations. The routing automatically redirects affected flows, but their packets are discarded by policers since they do not have a reservation on their deviation path. Protocol mechanisms may assure that a new reservation is set up, however, this cannot avoid service interruptions and the simultaneous requests of a multitude of redirected flows imposes an unrealistically high signalling load on routers. In contrast, resilient NAC reserves backup capacities in advance to protect rerouted traffic and to avoid heavy reservation signalling due to a transient network failure. The simplest and most efficient resilient NAC implementation is the enhancement of the BBB NAC [21] as it has a stateless core. The virtual capacity budgets $BBB(v, w)$ are just set low enough such that the redirection of admitted traffic cannot cause overload on any link when a failure occurs. The configuration of the budgets for resilient BBB NAC is well feasible and leads to reduced but still acceptable resource utilization. When we refer to AC in the following, we have in mind the non-resilient and resilient version of BBB NAC since we assume network resilience as a mandatory requirement for carrier grade networks.

2.4 Comparisons of AC and CO

We briefly address other comparisons of AC and CO to distinguish them from our work.

The work in [4] considers different utility functions for rigid and adaptive applications and different flow level models including the Poisson model. They are used to compare the additional capacity above the mean rate that is required for networks with reservations and for networks with a best effort service. In case of the Poisson model, they find only marginal benefits for AC vs. CO even for rigid applications while for other flow level models AC reveals clear benefits. The study pertains only to a single link such that questions like the impact of traffic shifts and redirected traffic are out of scope.

A comparison of AC and CO in access network dimensioning is the topic of [40]. They consider the aggregation link in a hierarchically structured access network and find a clear benefit of AC. Depending on parameters like blocking probability, packet loss probability, and user activity, the number of subscribers for a given access network capacity can be substantially higher when AC is used. In contrast, our work focuses on the dimensioning of an entire network and considers potential traffic shifts and redirected traffic.

The authors of [20] have shown that if the Poisson model is used for the characterization of the dynamics on the flow level, the required capacity for AC and CO is almost the same since the traffic variability of the Poisson model is rather low for highly aggregated traffic. Therefore, they developed the concept of single hot spots to model traffic shifts. Depending on the strength of the hot spots, AC leads to significant capacity savings compared to CO. In this work, we introduce general multi-hot-spots. Note that the work in [20] did not address resilience issues which constitute the main contribution of this study.

3. CAPACITY COMPARISON FOR CO AND AC ON A SINGLE LINK

In this section we present the basics for our performance analysis together with a sensitivity analysis regarding basic parameters. The objective is twofold. On the one hand, we want to provide a good understanding of the dimensioning methods since they are applied in Section 4 in a more complex context. On the other hand, we show that the input parameters for capacity dimensioning have a visible but moderate influence on the required capacity. As a consequence, their special choice has no impact on our general findings in Section 4.

3.1 Performance Analysis

We first explain our basic assumptions for the comparison of the required capacity for CO and AC and then we present the applied traffic model together with the capacity dimensioning algorithms for CO and AC.

3.1.1 Methodology of the Comparison

Above, we mentioned effective bandwidths which are allocated to flow requests in networks with AC to meet the required QoS. The effective bandwidth depends on the queuing behavior of the underlying packet level model. As we focus on the comparison regarding the resource efficiency in networks with CO and AC, we assume the same packet level model in both network types, which leads to the same required bandwidth for a flow. An inadequate packet level model leads to QoS degradation in both systems. In networks with AC, effective bandwidths for flows are certainly easier to determine than in networks with CO because flow descriptors provide helpful information and control mechanisms can enforce them. This is, however, not the concern of this comparison. We rather use this consideration to eliminate the uncertainty of the packet level model to facilitate a comparison of the resource requirements of CO and AC.

3.1.2 Model for a Traffic Aggregate

We model the number of active flows of a traffic aggregate and present our assumption on their effective bandwidths.

3.1.2.1 Flow Generation according to the Poisson Model.

We consider networks with realtime flows. Such a setting may be found in a dedicated network for realtime traffic like the UMTS core network or in the Differentiated Services architecture when we focus only on the bandwidth for high priority traffic. The Poisson model for flow arrivals is appropriate for Internet traffic [5, 6, 27, 30] and current evidence of Poisson inter-arrival times for VoIP calls is given in [10]. Therefore, we use a flow level model that is characterized by an exponentially distributed inter-arrival time and a general, independently and identically distributed call holding time. The offered load of a system is its average number of simultaneous flows if no flow blocking occurs due to AC. It is measured in the pseudo unit ‘‘Erlang’’ and it is calculated by $a = \frac{\lambda}{\mu}$ where λ is the arrival rate and $\frac{1}{\mu}$ the mean holding time of the flows.

3.1.2.2 Traffic Mix.

We use the simplified multirate model from [21] with $n_r = 3$ different rate types r_0 , r_1 , and r_2 with a bit rate of $c(r_0) =$

64 kbit/s, $c(r_1) = 256$ kbit/s, and $c(r_2) = 2048$ kbit/s. The random variable R_{t_R} indicates the requested rate in case of a flow arrival. Its distribution in Table 1 is parameterized such that the mean rate of the flows $E(c(R_{t_R})) = 256$ kbit/s is independent of the parameter $t_R \in [0, 1]$ and that the coefficient of variation of their rate $c_{var}(c(R_{t_R})) = 2.291 \cdot t_R$ depends linearly on it. We assume that the flows of all rate types have the same mean holding time. Therefore, we can calculate the rate-specific offered load by $a_i = a \cdot p(r_i)$.

Table 1: The distribution of the flow rate R_{t_R} (effective bandwidth) depends on the parameter $t_R \in [0, 1]$.

| request type r_i | $c(r_i)$ | $p(r_i)$ |
|--------------------|-------------|-----------------------------|
| r_0 | 64 kbit/s | $\frac{28}{31} \cdot t_R^2$ |
| r_1 | 256 kbit/s | $(1 - t_R^2)$ |
| r_2 | 2048 kbit/s | $\frac{3}{31} \cdot t_R^2$ |

3.1.3 Capacity Dimensioning for AC Using the Multirate $M/G/n - 0$ Queue

For the sake of simplicity, we explain capacity dimensioning for AC prior to the one for CO. AC limits the number of flows to prevent overload. It blocks a new flow if its effective bandwidth together with the sum of the effective bandwidth of the admitted flows exceeds the link bandwidth. The probability for a flow to be blocked at its arrival is denoted by $p_b^i(C)$ which depends on its flow rate r_i and the link capacity C . We derive this flow blocking probability from a multirate $M/G/n - 0$ queue without buffers. The request size of the flows is an integral multiple of a basic capacity unit $u_c = 64$ kbit/s for the above rate distribution, i.e., each request with rate $c(r_i)$ can be expressed as an integral multiple of that unit by $c_u(r_i) = \frac{c(r_i)}{u_c}$. The number of servers n of the queue is the link capacity C measured in capacity units. The active flows together with their request sizes determine the used bandwidth which corresponds – expressed in capacity units u_c – to the number of busy servers. This corresponds to the system state, i.e. to the number of customers, of the $M/G/n - 0$ queue. The state changes due to the stochastic behavior of the multirate Poisson model and the state probabilities can be calculated by the well-known Kaufman/Roberts algorithm presented in [29].

Blocking is experienced by a newly arriving flow f when the system is in a state with insufficient free capacity to accommodate its request size $c(r(f))$. The blocking probability for f is the sum of the probabilities of all states in which blocking occurs for a flow with rate $c(r(f))$. This leads to the observation that flows with large request rates face a larger blocking probability. The blocking probability $p_b(C)$ we use for capacity dimensioning in our study is the average of the blocking probabilities $p_b^i(C)$ of all request types r_i , weighted by their occurrence probability $p(r_i)$ and rate $c(r_i)$. Thus, we can dimension the link capacity $C = n \cdot u_c$ by choosing the number of servers n large enough that admission requests are rejected only with a very small target blocking probability p_b : $C = \min_{C'} \{p_b(C') \leq p_b\}$. The algorithm in [21] calculates this number in an efficient way.

3.1.4 Capacity Dimensioning for CO Using the Multirate $M/G/\infty$ Queue

We adapt the above model to CO. As CO does not block any flows, the number of flows on the link is not bounded. Therefore, we model a link by a multirate $M/G/\infty$ queue

with infinitely many server units to avoid blocked requests. The calculation of the equilibrium state probabilities of the number of busy servers is known as the stochastic knapsack with infinite capacity [31]. The state equilibrium can be calculated as follows. The request types constitute $k = n_r$ classes and the k -dimensional state space is described by $\mathcal{X} = \{x = (x_0, x_1, \dots, x_{k-1}) \in \mathbb{N}_0^k\}$ whereby the x_i denote the number of flows of request type r_i in the system. The equilibrium state probabilities are

$$p(x) = \prod_{i=0}^{k-1} \frac{a_i^{x_i}}{x_i!} e^{-a_i} \quad (1)$$

with a_i being the class-specific offered load in Erlang. The consideration of the type-specific rates $c(r_i)$ yields the required link capacity $c(x) = \sum_{i=0}^{k-1} c(r_i) \cdot x_i$ of state x .

We discuss two different QoS violation probabilities for CO that depend both on the link bandwidth C .

p_v^f The first definition is consistent with the definition of the flow blocking probability p_b . It is the QoS violation seen by a newly arriving flow f . This probability p_v^f comprises the probability of all states in which a new flow sees a QoS violation after its arrival.

$$p_v^f(C) = 1 - \sum_{0 \leq i < n_r} p(r_i) \cdot \sum_{\{x \in \mathcal{X}: c(x) < C - c(r_i)\}} p(x) \quad (2)$$

p_v^t The second definition is the QoS violation probability p_v^t over time. Thus, it is simply

$$p_v^t(C) = 1 - \sum_{\{x \in \mathcal{X}: c(x) \leq C\}} p(x) \quad (3)$$

Note that probability $p_v^t(C)$ is smaller than $p_v^f(C)$.

An overprovisioned link requires so many capacity units C that the probability for the flows to exceed this bandwidth is smaller than a given tolerable violation probability $p_v^{\{f,t\}}$. Thus, the required capacity is

$$C = \min_{C'} \{p_v^{\{f,t\}}(C') \leq p_v^{\{f,t\}}\}. \quad (4)$$

3.2 Numerical Results

We illustrate the above dimensioning methods by studying the impact of various parameters on the required link capacity. We also assess the missing capacity for overprovisioned links and argue for an enhancement of the traffic model that leads to more overload than the Poisson model.

3.2.1 Impact of the Dimensioning Method on the Required Capacity

We dimension a single link for CO and AC with $p_v^f, p_v^t, p_b = 10^{-3}$ for a single effective bandwidth of 256 kbit/s per flow, i.e. $t_R = 0$ (cf. Table 1). Figure 1 shows the absolute required capacity depending on the offered link load given in Erlang. The jerky curves at the left margin of the figure are due to granularity effects for small offered load. Apart from that, the absolute required capacity increases almost linearly with the offered load for more than 100 Erlang, but the lines hardly differ and it is hard to interpret the results. Therefore, we also plot the relative required capacity as a multiple of the average offered traffic in the same figure. It clearly shows that the relative amount of additional capacity

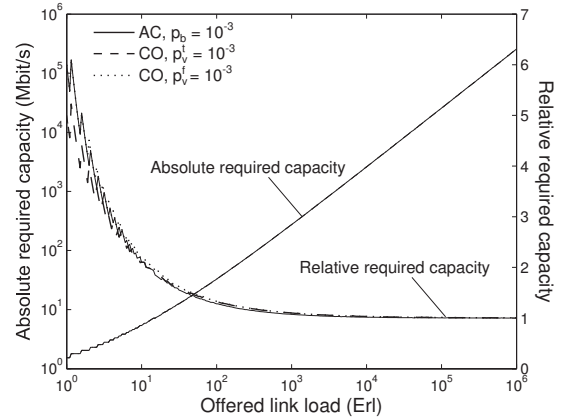


Figure 1: Impact of the offered load on the absolute and relative required capacity of a single link for various dimensioning methods.

decreases with an increasing offered load. This fact is called economy of scale.

If p_v^f is used for capacity dimensioning for CO, AC requires less capacity than CO since AC blocks some of the traffic and reduces thereby slightly the load in the system compared to CO. However, if p_v^t is applied as capacity dimensioning objective, the required capacity for CO is reduced to such an extent that it is smaller than the one for AC for very small offered load. However, the difference between all these methods is negligible for medium or large offered load. In the following, we denote p_v^t simply by p_v and use it as the objective for capacity dimensioning with CO since QoS violation caused by a single flow hits all flows in progress and not just the new one. This is unlike with AC where only the arriving flow is blocked.

3.2.2 Impact of the Request Rate Variance and the Target Probabilities on the Capacity

We investigate the impact of the objective probabilities p_v and p_b and the parameter t_R for the distribution of the effective flow bandwidth (cf. Table 1) on the required link capacity. We consider both a homogeneous traffic mix ($t_R = 0$) and a strongly heterogeneous traffic mix ($t_R = 1$) for the objective probabilities $p_v, p_b = 10^{-3}$ and 10^{-6} . The results are compiled in Figure 2. For $p_v = p_b$, AC and CO need almost the same amount of capacity. Smaller objective probabilities and more heterogeneous effective bandwidth increase the required link capacity significantly, but only for little offered load. The influence of the variability of the effective bandwidth is clearly stronger than the one of the target probabilities. In the following we use $t_R = 1$ since it is more realistic than $t_R = 0$ for Internet flows whose request rates can be highly variable.

3.2.3 Impact of the Target Probability for CO on the QoS Violation

The QoS violation perceived by the user depends on the severity of the missing capacity. Therefore, we calculate the average of the missing capacity in case of CO relative to the provisioned capacity C by

$$E[M] = \frac{1}{C} \cdot \sum_{\{x \in \mathcal{X}: c(x) > C\}} p(x) \cdot (c(x) - C) \quad (5)$$

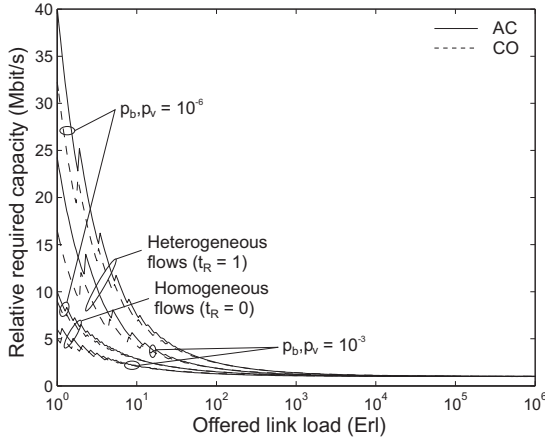


Figure 2: Impact of the distribution of the effective flow bandwidth and the objective probabilities for capacity dimensioning on the required capacity of a single link.

where x is the state vector of flows in the system. The missing capacity is evaluated for the above experiments with heterogeneous traffic for $p_v = 10^{-3}$ and 10^{-6} and the results are shown in Figure 3 in percent. The jerkiness of the graph is again caused by granularity effects. The scaling of the left y-axis immediately shows that the average of the missing capacity in percent is in the order of p_v , i.e. 10^{-3} and 10^{-6} , respectively. We also calculate the conditional average of the missing capacity for overload situations by $E[M_c] = \frac{E[M]}{p_v}$. According to the construction of the graph, the curves for $E[M]$ and $E[M_c]$ coincide, but they adhere to different y-axes. When the QoS is violated, approximately 4% or 8% capacity is missing for little offered load and for a target probability of $p_v = 10^{-6}$ and 10^{-3} , respectively. Medium offered load misses around 1% or 2%, and for large offered load the missing capacity is negligible regardless of the target probability p_v . These values are surprisingly low which results from the smooth behavior of the Poisson model and the fact that we assumed constant offered load in our experiment. This allows only small statistical oscillations and does not model overload due to increased content attractiveness at certain locations.

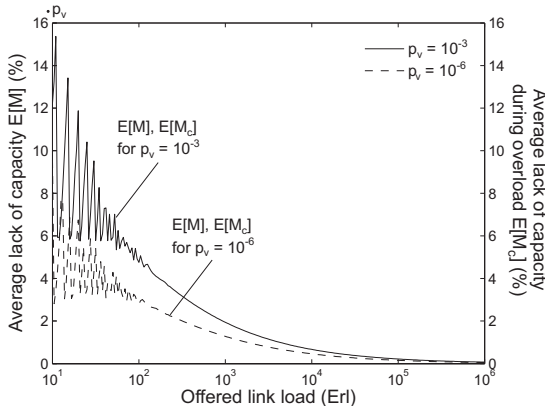


Figure 3: Impact of the offered load and the target probability p_v on the overall and conditional average QoS violation $E[M]$ and $E[M_c]$ for CO.

For the rest of the paper we choose a maximum flow blocking probability of $p_b = 10^{-3}$ for AC and a maximum QoS violation probability of $p_v = 10^{-6}$ for CO. The difference is motivated by the fact that flow blocking is annoying for the affected user, but QoS violation hits all flows in progress and should be avoided. Note that the required capacity and the QoS violation revealed only little sensitivity to these parameters for medium and large offered load. The required capacity is mainly controlled by the offered load and, therefore, the choice of the target probability has only a minor impact on the results in Section 4.

3.2.4 Impact of Transient Overload on the Capacity

We assume a constant offered load for most of the time and a temporary increase of the normal offered load by an overload factor of f_l . AC can block excess traffic during time of overload and preserve QoS at the expense of blocked flows. In contrast, CO must provide so much capacity that the excess traffic can be carried. Figure 4 shows the required capacity for CO and AC together with the flow blocking probability p_b^o for AC during time of overload. The results are shown for an offered load of $a = 10^2$ and 10^5 Erlang in the non-overload case. As the required capacity for AC is dimensioned for the non-overload case, the respective curves are independent of the overload factor. However, the blocking probability for AC increases with the overload factor f_l . The blocking probability for 10^5 Erlang is larger than the one for 10^2 Erlang because there is less additional capacity available relative to the average traffic rate due to economy of scale. The overload factor $f_l = 1$ denotes the non-overload case for CO. CO requires visibly more capacity than AC for $a = 10^2$ Erlang because it uses $p_v = 10^{-6}$ as target probability for dimensioning instead of $p_b = 10^{-3}$. However, for $a = 10^5$ Erlang, the capacity requirements for CO and AC are almost equal for $f_l = 1$. With an increasing overload factor f_l , the required capacity for CO is scaled up about linearly since it must be dimensioned for the offered load during the overload interval. In fact, this result for a single link is trivial. Therefore, in the next section, we consider in networks different sources of overload that do not increase the overall traffic. We rather model traffic shifts that cause local overload.

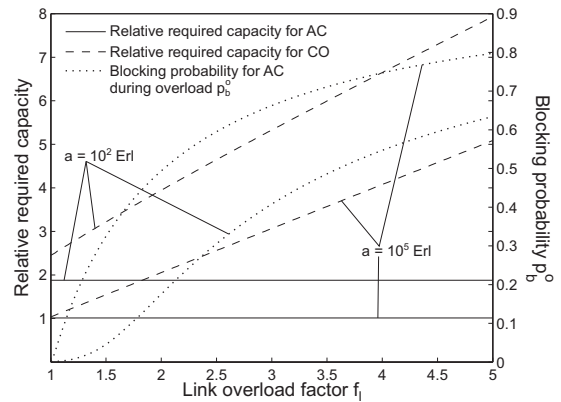


Figure 4: Impact of the overload factor f_l on the required capacity and the blocked traffic of a single link for $a = 10^2$ and 10^5 Erlang.

4. CAPACITY COMPARISON FOR CO AND AC IN NETWORKS

The Poisson model accounts for the stochastic fluctuations of the number of flows in the traffic aggregate. However, if the offered load is constant, it produces very smooth traffic rates such that only little additional capacity is needed both in networks with CO and AC. In this section, we investigate the impact of overload that results from traffic shifts within the network which temporarily increases the offered load on individual links without increasing the overall traffic in the network. Such traffic shifts may result from increased content attractiveness at certain locations or from redirected traffic due to network failures. For both issues we need to consider the entire network instead of a single link. We first extend our performance analysis to networks and investigate then the impact of traffic shifts and redirected traffic on the required capacity for CO and AC.

4.1 Extension of the Performance Analysis to Networks

We extend the traffic model and the dimensioning methods for CO and AC from Section 3.1.2 to networks. Finally, we adapt the definition of the relative required capacity from a single link and introduce a compact description of the networking scenarios we consider in the next section.

4.1.1 Extension of the Traffic Model

We review a simple method for the generation of a basic traffic matrix from which we derive traffic matrices with simple and complex traffic shifts [20]. Finally, we introduce a notation for network failures and the resulting (re)routing that also leads to traffic shifts.

4.1.1.1 Generation of the Basic Traffic Matrix.

Most of the network experiments in this paper are based on the Labnet03 reference network given in Figure 5 from [21]. We denote the topology of a network by a set of nodes \mathcal{V} and a set of bidirectional edges \mathcal{E} . All network nodes are both ingress and egress routers. The average border-to-border (b2b) load between two nodes in the network is denoted by a_{b2b} . It determines the overall offered load in the network $a_{tot} = \sum_{v,w \in \mathcal{V}, v \neq w} a(v,w) = |\mathcal{V}| \cdot (|\mathcal{V}|-1) \cdot a_{b2b}$ whereby $|\mathcal{V}|$ denotes the cardinality of the set \mathcal{V} . The generation of the traffic matrix is based on the population of the cities and their surroundings that are compiled in [21]. For two cities corresponding to the nodes v and w with population sizes $\pi(v)$ and $\pi(w)$, the b2b offered load $a(v,w)$ amounts to

$$a(v,w) = \begin{cases} \frac{a_{tot} \cdot \pi(v) \cdot \pi(w)}{\sum_{x,y \in \mathcal{V}, x \neq y} \pi(x) \cdot \pi(y)} & \text{for } v \neq w, \\ 0 & \text{for } v = w. \end{cases} \quad (6)$$

4.1.1.2 Hot Spot Model Causing Transient Traffic Shifts.

We model a hot spot in the network by increasing the traffic attraction of a single city v by a hot spot factor f_h that is expressed by a modified population function

$$\pi_{overload}^v(w) = \begin{cases} \pi(w) & \text{if } w \neq v \\ f_h \cdot \pi(w) & \text{if } w = v \end{cases} \quad (7)$$

which is used as input for Equation (6). This overload model

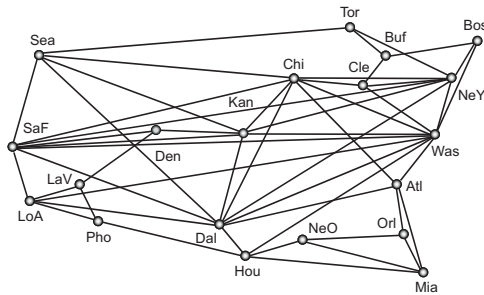


Figure 5: Topology of the Labnet03 network with 20 nodes and 53 bidirectional links.

is quite conservative since it does not increase the overall traffic in the network. It just causes a traffic shift which changes only the structure of the traffic matrix. As a consequence, an increased or decreased load on individual links can be observed. Note that every node $v \in \mathcal{V}$ is a potential hot spot and even several hot spots may occur simultaneously. Therefore, we characterize a hot spot scenario by the set of routers with increased attractiveness, e.g. $h = \{v, w\}$. In the following, \mathcal{H} denotes the set of considered hot spot scenarios and it contains always the normal scenario $h = \emptyset$. Note that traffic variation may also be caused by other influences, e.g. inter-domain rerouting [32], and they may increase the overall traffic volume in the network.

4.1.1.3 Description of Network Failures and Routing.

The connectivity of the network after a failure depends on the network topology and the applied restoration or protection switching mechanism. In our experiments, we use shortest path routing since it is the basis for the most frequently used Interior Gateway Protocols (IGPs) such as OSPF and IS-IS. We characterize a network failure s by the set of failed network elements, e.g. links or routers. In our investigation, we consider only single link failures for which the networks under study remain fully connected after rerouting. The resilience against network failures regarding QoS depends both on the connectivity in a failure scenario and on the available capacity on the backup paths. Thus, we dimension the required capacity in such a way that it prevents overload due to the redirected traffic for a limited set of protected failure scenarios \mathcal{S} . This set contains the failure-free case $s = \emptyset$ by default.

The traffic aggregate between v and w is denoted by $g(v,w)$ and the set of all aggregates in the network is \mathcal{G} . The routing of an aggregate $g \in \mathcal{G}$ within the network depends on the failure case s . We describe it by the function $u(s,l,g)$ that describes the percentage of the traffic rate $c(g)$ that uses link l in a specific failure case $s \in \mathcal{S}$, i.e., the routing in the failure-free case is given by the function $u(\emptyset,l,g)$. This notation is very general since it can express the routing of arbitrary restoration and protection mechanisms and copes well with load balancing.

4.1.2 Extension of the Capacity Dimensioning Methods to Networks in the Presence of Traffic Shifts and Network Failures

We extend the capacity dimensioning methods for a single link from Section 3.1.3 and 3.1.4 to networks and adapt them to traffic shifts and network failures. This constitutes the

concept of “resilient capacity overprovisioning” which is part of the the main contribution of this paper.

A concrete networking scenario $z = (h, s)$ is determined by its traffic matrix which depends on the hot spot scenario h , and the failure scenario s . Conversely, the functions $h(z)$ and $s(z)$ yield the respective hot spot and failure scenarios. In the following, we explain the dimensioning methods for CO and for AC in networks starting with a single specific networking scenario z and extending it to a set of considered networking scenarios \mathcal{Z}_c .

4.1.2.1 Dimensioning of Link Capacities in a Network for CO.

The offered load $a(z, l)$ for the link l in a specific networking scenario z can be calculated by taking into account the load contribution of all traffic aggregates for that link:

$$a(z, l) = \sum_{g \in \mathcal{G}} a(h(z), g) \cdot u(s(z), l, g). \quad (8)$$

Based on this value and a suitable target probability p_v , the capacity dimensioning algorithm for CO presented in Section 3.1.4 computes the capacity $c(z, l)$ of that link for the networking scenario z . The required link capacity for a set of considered networking scenarios \mathcal{Z}_c is then simply the maximum link capacity of all its networking scenarios $z = \mathcal{Z}_c$:

$$c(l) = c(\mathcal{Z}_c, l) = \max_{z \in \mathcal{Z}_c} (c(z, l)). \quad (9)$$

4.1.2.2 Dimensioning of Link Capacities in a Network for AC.

We dimension the capacity for the BBB NAC described in Section 2.3.2 since this NAC method is resilient to network failures if it is configured appropriately. For each traffic aggregate $g \in G$ a b2b budget exists with a capacity of $c(g)$ that can be dimensioned based on the offered load $a(\emptyset, g)$ with the link dimensioning algorithm for AC presented in Section 3.1.3. Note that in networks with resilient AC, failures but no hot spots need to be respected since overload due to hot spots can be blocked. Thus, the capacity for link l in the networking scenario z can be determined by

$$c(z, l) = \sum_{g \in \mathcal{G}} c(g) \cdot u(s(z), l, g) \quad (10)$$

and the required capacity for a set of considered networking scenarios \mathcal{Z}_c is again calculated according to Equation (9).

4.1.3 Performance Measure and Selected Sets of Networking Scenarios

For the sake of easier notation in the next section, we shortly describe the performance measure for networks and some selected sets of networking scenarios.

4.1.3.1 Performance Measure.

Like in Section 3.2, we use again the relative required capacity as performance measure. However, we need to adapt its definition from a single link to a network. The absolute required network capacity is $C_{abs} = \sum_{l \in \mathcal{E}} c(l)$. The average traffic rate under normal conditions can be calculated by $C_{avg} = E(c(R_{tR})) \cdot \sum_{l \in \mathcal{E}} a(z = (\emptyset, \emptyset), l)$. Thus, we define the relative required network capacity by $C_{rel} = \frac{C_{abs}}{C_{avg}}$.

4.1.3.2 Selected Sets of Networking Scenarios.

We define sets of networking scenarios that are of particular interest for our study. We assess their size for our test network in Figure 5 to give an idea of the complexity of the investigation. The sets $\mathcal{Z}^{i,0}$ contain only failure-free networking scenarios for the investigation of the impact of single and double hot spot scenarios without link failures.

- $\mathcal{Z}^{0,0} = \{(\emptyset, \emptyset)\}$; “the basic traffic matrix in the failure-free scenario”, $|\mathcal{Z}^{0,0}| = 1$.
- $\mathcal{Z}^{1,0} = \mathcal{Z}^{0,0} \cup \{ \text{“all single hot spots in the failure-free scenario”} \}$, $|\mathcal{Z}^{1,0}| = |\mathcal{Z}^{0,0}| + \binom{|\mathcal{E}|}{1} = 1 + 20 = 21$.
- $\mathcal{Z}^{2,0} = \mathcal{Z}^{1,0} \cup \{ \text{“all double hot spots in the failure-free scenario”} \}$, $|\mathcal{Z}^{2,0}| = |\mathcal{Z}^{1,0}| + \binom{|\mathcal{V}|}{2} = 21 + 190 = 211$.

The sets $\mathcal{Z}^{i,1}$ contain networking scenarios with all single link failures for the investigation of the impact of single and double hot spot scenarios in the presence of link failures.

- $\mathcal{Z}^{0,1} = \mathcal{Z}^{0,0} \cup \{ \text{“all single link failure scenarios without hot spots”} \}$, $|\mathcal{Z}^{0,1}| = |\mathcal{Z}^{0,0}| + \binom{|\mathcal{E}|}{1} = 1 + 53 = 54$.
- $\mathcal{Z}^{1,1} = \mathcal{Z}^{0,1} \cup \{ \text{“all single link failure scenarios combined with all simultaneous single hot spots”} \}$, $|\mathcal{Z}^{1,1}| = |\mathcal{Z}^{0,1}| + |\mathcal{Z}^{0,1}| \cdot \binom{|\mathcal{V}|}{1} = 54 + 54 \cdot 20 = 1134$.
- $\mathcal{Z}^{2,1} = \mathcal{Z}^{1,1} \cup \{ \text{“all single link failure scenarios combined with all simultaneous double hot spots”} \}$, $|\mathcal{Z}^{2,1}| = |\mathcal{Z}^{1,1}| + |\mathcal{Z}^{0,1}| \cdot \binom{|\mathcal{V}|}{2} = 1134 + 54 \cdot 190 = 11394$.

4.2 Numerical Results

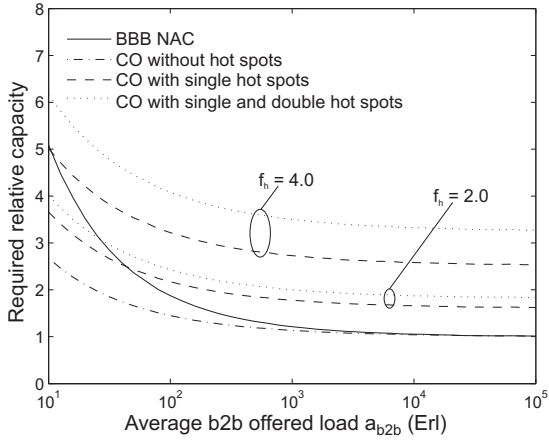
We study the required overall capacity for networks with CO in the presence of simple and complex traffic shifts, as well as with and without resilience requirements. We compare the results with those for networks with AC. We conduct the comparisons both in the Labnet03 network (cf. Figure 5) and in random networks of different size.

4.2.1 Capacity Requirements in Non-Resilient Networks

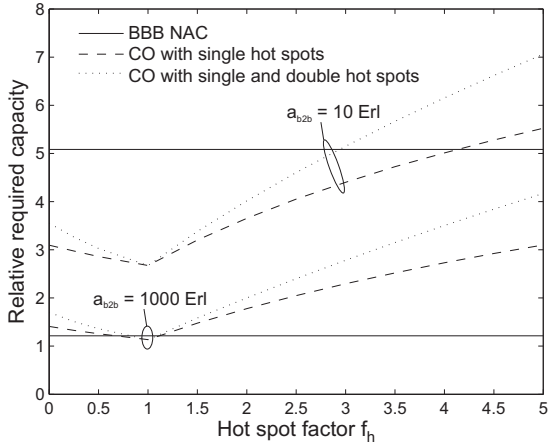
We illustrate the impact of hot spot scenarios on the required capacity for CO and AC in non-resilient networks.

4.2.1.1 Experiments with Labnet03.

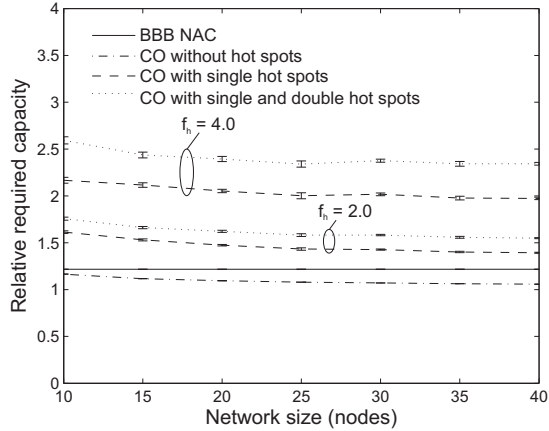
Figure 6(a) shows the relative required capacity in the Labnet03 network depending on the average offered load a_{b2b} between two border routers. The network capacity is dimensioned for the considered networking scenarios $\mathcal{Z}^{0,0}$ (without hot spots), $\mathcal{Z}^{1,0}$ (single hot spots only), and $\mathcal{Z}^{2,0}$ (single and double hot spots), and for BBB NAC. The hot spot factor is set to $f_h = 2$ and to $f_h = 4$, respectively. Like in the single link experiments, the relative required capacity decreases for all curves with an increasing load. Surprisingly, CO without hot spots ($\mathcal{Z}^{0,0}$) requires less capacity than AC. The reason is that CO can take advantage of the fact that the offered load on a link is larger than the load for a single budget. The capacity dimensioning for a specific link for CO is based on the overall load of all aggregates carried over this link (cf. Equation (8)) while the BBB NAC considers only the load of a single aggregate for a b2b budget. Thus, in contrast to AC, CO benefits from increased economy of scale which leads to less required capacity for CO than for



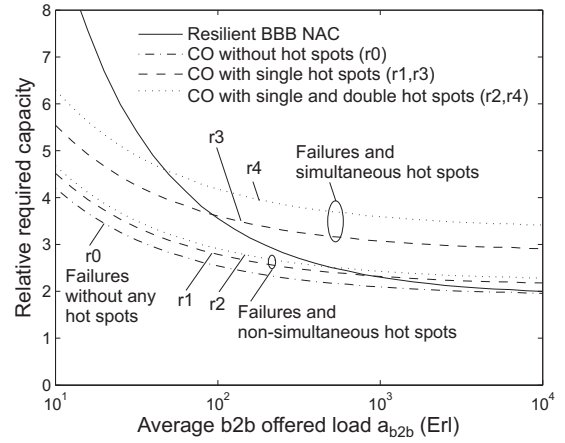
(a) Influence of the b2b offered load a_{b2b} in Labnet03 for hot spot factors of $f_h = 2$ and 4.




(b) Influence of the hot spot factor f_h in Labnet03 for a b2b offered load of $a_{b2b} = 10$ and 1000 Erlang.



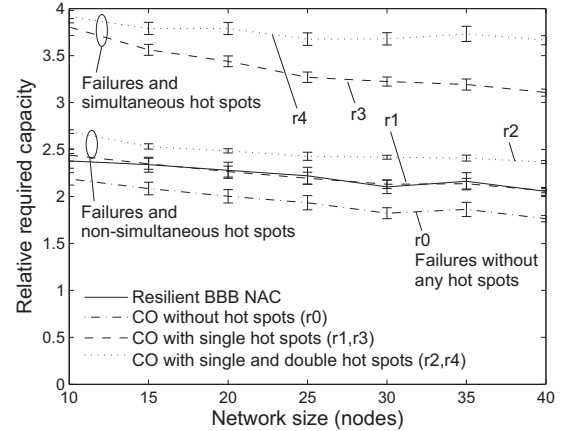
(c) Influence of the network size for a b2b offered load of $a_{b2b} = 1000$ Erlang and hot spot factors $f_h \in \{2, 4\}$.



(a) Influence of the b2b offered load a_{b2b} in Labnet03 for a hot spot factor of $f_h = 2$.



(b) Influence of the hot spot factor f_h in Labnet03 for a b2b offered load of $a_{b2b} = 1000$ Erlang.



(c) Influence of the network size for a b2b offered load of $a_{b2b} = 1000$ Erlang and a hot spot factor of $f_h = 2$.

Figure 6: Relative required capacity in non-resilient networks with capacity overprovisioning (CO) and admission control (AC), respectively.

Figure 7: Relative required capacity in resilient networks with capacity overprovisioning (CO) and admission control (AC), respectively.

AC. For sufficiently large offered load, AC works efficiently, too. CO with single hot spots requires more capacity than AC since it must provide enough resources for all possible traffic shifts. CO for double hot spots needs visibly more resources than CO for single hot spots. An increase of the hot spot factor from $f_h = 2$ to $f_h = 4$ also increases the resource requirements for CO considerably.

Figure 6(b) shows the relative required network capacity for an offered b2b load of $a_{b2b} = 10$ and 1000 Erlang depending on the hot spot factor f_h . The capacity curves for $a_{b2b} = 1000$ Erlang reveal an almost linear growth which is smaller than f_h . This is different to the experiment on the single link (cf. Figure 4) which can be explained as follows. The links adjacent to a hot spot carry all the “hot spot traffic” from and to this hot spot. The rate of these aggregates scales almost with f_h . However, the transit traffic on these links is hardly affected or even decreased by the hot spot. As a consequence, the required capacity for the adjacent links grows less than by f_h since their carried traffic consists of increased hot spot and slightly decreased transit traffic. The capacity curves for single hot spots require less resources than those for double hot spots. They meet for $f_h = 1$ since this is the value for CO without any hot spots. Hot spot factors $f_h < 1$ produce “cold spots”, i.e., the attractiveness of a certain node is reduced which also effects a traffic shift. However, a cold spot leads only to a small increase of the required capacity. The required network capacity for AC is independent of the hot spot factor and produces, therefore, horizontal lines. For very little offered load of $a_{b2b} = 10$ Erlang, AC requires significantly more resources than CO, but for a large offered load of $a_{b2b} = 1000$ Erlang, AC works efficiently enough such that it can effectively save capacity by blocking excess traffic in overload situations.

4.2.1.2 Experiments with Random Networks.

We investigate the impact of the network size on the relative required network capacity for CO and AC. To that end, we construct random networks with n nodes and an average node degree of $deg_{avg} = 3$, i.e. with $m = \frac{n \cdot deg_{avg}}{2}$ bidirectional links using the algorithm given in [21]. This algorithm guarantees a connected graph and keeps the degree of every node between $2 \leq deg_{avg} \leq 4$. Like above, we dimension the capacity of these networks for CO without hot spots, with single hot spots only, and with single and double hot spots. The results are illustrated in Figure 6(c) for an average b2b offered load of $a_{b2b} = 1000$ Erlang and for hot spot factors of $f_h = 2$ and 4. The relative required capacity for CO without hot spots decreases slightly for an increasing network size since larger networks lead to more offered load per link and thereby to increased economy of scale for CO. In contrast, BBB NAC cannot benefit from that. For a hot spot factor of $f_h = 2$, single hot spots only lead to about 50% more capacity whereas single and double hot spots lead to 75% more capacity than the average traffic rate in the network. Doubling the hot spot factor to $f_h = 4$ also doubles the additional capacity requirements to 100-150%.

4.2.2 Capacity Requirements in Resilient Networks

We illustrate the impact of hot spot scenarios on the required capacity for CO and AC in networks with resilience requirements.

4.2.2.1 Experiments with Labnet03.

We consider CO for the following 5 types of networking scenarios.

- (r0) $\mathcal{Z}_c = \mathcal{Z}^{0,1}$, i.e. resilience against link failures without elasticity for any hot spots.
- (r1) $\mathcal{Z}_c = \mathcal{Z}^{0,1} \cup \mathcal{Z}^{1,0}$, i.e. resilience against link failures with elasticity for non-simultaneous single hot spots.
- (r2) $\mathcal{Z}_c = \mathcal{Z}^{0,1} \cup \mathcal{Z}^{2,0}$, i.e. resilience against link failures with elasticity for non-simultaneous single and double hot spots.
- (r3) $\mathcal{Z}_c = \mathcal{Z}^{1,1} \cup \mathcal{Z}^{2,0}$, i.e. resilience against link failures with elasticity for non-simultaneous single and double hot spots and simultaneous single hot spots.
- (r4) $\mathcal{Z}_c = \mathcal{Z}^{2,1} \cup \mathcal{Z}^{2,0}$, i.e. resilience against link failures with elasticity for simultaneous and non-simultaneous single and double hot spots.

To assess the relevance of these networking scenarios for capacity dimensioning in practice, we make the following considerations. We assume the probability of a link failure to be smaller than the one for a hot spot, i.e. $p_l < p_h$. Single link failures must be protected as well as double hot spots. However, we expect that the simultaneous occurrence of a single link failure together with a simultaneous hot spot is so unlikely that we do not need to provide sufficient capacity for those scenarios. Under these assumptions, option r2 is appropriate for resilient CO in practice.

Figures 7(a) and 7(b) show the relative required capacity for CO and AC with resilience against single link failures in Labnet03. They correspond to Figures 6(a) and 6(b), but we show the results for the above mentioned option only for $f_h = 2$. Figure 7(a) shows that resilient CO and AC require both substantially more capacity than CO or AC without resilience against link failures since they both require backup capacities for redirected traffic on the links. The limit for large a_{b2b} depends on the network topology and the applied restoration or protection switching mechanism. Note that the backup capacity can be minimized by routing optimization [23]. The curves for resilient CO (r1) and (r2) require only marginally more capacity than the curve for (r0). This means that the backup capacity for single link failures almost suffices to absorb traffic shifts due to single and double hot spots for a hot spot factor of $f_h = 2$. As a consequence, resilient CO for application in practice (r2) requires only little more capacity than resilient BBB NAC. We also plotted the options r3 and r4 for resilient CO in the figures to illustrate that they need about 100% more capacity than r0, r1, and r2. This extra capacity is needed to accommodate redirected traffic caused by link failures during simultaneous single or even double hot spots.

Figure 7(b) keeps the offered load fixed at $a_{b2b} = 1000$ Erlang and varies the hot spot factor f_h . Resilient CO (r2) is as efficient as resilient AC for hot spot factors up to about $f_h = 2$. The two figures show that the relative required capacity for resilient CO depends on the offered load a_{b2b} , the hot spot factor f_h , and the resilience option. In contrast, for resilient AC it depends only on the offered load a_{b2b} .

4.2.2.2 Experiments with Random Networks.

Finally, we apply the above option to the random networks from Section 4.2.1 such that the results in Figure 7(c)

are comparable to those in Figure 6(c). The figure shows that resilient CO without elasticity for simultaneous hot spots during link failures (r0, r1, r2) needs a similar amount of capacity like resilient AC for $a_{b2b} = 1000$ and $f_h = 2$. Resilient CO with elasticity for simultaneous hot spots requires again about 100% more additional resources. This observation is apparently independent of the network size.

5. CONCLUSION

Capacity overprovisioning (CO) must provide sufficient capacity on the links that QoS degradation due to overload in the network is avoided. CO takes into account any kind of overload: (a) overload due to statistical variations of the normal traffic matrix, (b) overload due to changed traffic matrices caused by traffic shifts through popular sites (or by changes of the inter-AS routing), and (c) overload due to redirected traffic caused by network failures. This paper is the first to tackle all three sources of overload. It introduces the notion of resilient overprovisioning and proposes a capacity dimensioning method in such a way that the QoS violation probability p_v is kept below a given limit for the considered networking scenarios $z \in \mathcal{Z}_c$. This simple method is especially useful for a comparison of CO with AC methods. In addition, the idea of resilient CO can be certainly adapted to other traffic and overload models, e.g. to overload caused by routing changes of inter-AS routing.

Admission control (AC) is the counterpart to CO. We argued for resilient AC because the majority of overload situations in the Internet results from network failures [19]. We dimensioned the link capacities for networks with AC in such a way that the flow blocking probabilities p_b are kept low.

We assessed the impact of all three sources of overload on the required capacity by the “relative required capacity” which is the required capacity relative to the average traffic rate. We compared them for networks with CO and AC whereby the offered system load, the strength of traffic shifts, and the network size were key parameters for our investigation. The most important results of our study are the following.

- The target probabilities p_v and p_b for capacity dimensioning have only a small impact on the required capacity for CO and AC.
- The statistical fluctuations of the Poisson model for flows do not lead to significant overload and QoS violations. Therefore, additional overload models are needed.
- In networks without hot spots and failures, CO requires about the same capacity as AC or even less as it can take better advantage of economy of scale.
- Single hot spot scenarios lead to a significant increase of the required capacity for CO.
- Additional double hot spot scenarios increase these capacity requirements slightly.
- Resilience against link failures leads to increased capacity requirements for networks with CO and AC since both types require backup capacity for the redirected traffic.

- Resilient CO requires about the same network capacity as resilient AC to protect the network against failures and against overload due to single and double hot spots because the backup capacity can be used to absorb hot spots.
- We made these observations in a test network and confirmed them by a study of random networks of different size.

These findings can be generalized to other sources of overload, e.g. changes of the interdomain routing, since backup capacity can be reused to protect QoS against any kind of overload. Finally, we conclude that CO is even more attractive than AC in networks with resilience requirements.

6. ACKNOWLEDGMENTS

The authors would like to thank Prof. Tran-Gia for the stimulating environment which was a prerequisite for that work.

7. REFERENCES

- [1] A. Autenrieth and A. Kirstädter. Engineering End-to-End IP Resilience Using Resilience-Differentiated QoS. *IEEE Communications Magazine*, 40(1):50–57, Jan. 2002.
- [2] J. Beran, R. Sherman, M. Taqqu, and W. Willinger. Long-Range Dependence in Variable-Bit Rate Video Traffic. *IEEE Transactions on Communications*, 43(2/3/4):1566–1579, Feb. 1995.
- [3] N. Blefari-Melazzi and M. Femminella. Stateful vs. Stateless Admission Control: Which Can Be the Gap in Utilization Efficiency? In *IEEE Infocom*, Nov. 2002.
- [4] L. Breslau and S. Shenker. Best-Effort versus Reservations: A Simple Comparative Analysis. *ACM SIGCOMM Computer Communications Review*, 28:3–16, Sept. 1998.
- [5] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. On the Nonstationarity of Internet Traffic. In *ACM SIGMETRICS*, 2001.
- [6] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. *Nonlinear Estimation and Classification*, chapter Internet Traffic Tends Toward Poisson and Independent as the Load Increases. Springer, New York, 2002.
- [7] J. Cao and K. Ramanan. A Poisson Limit for Buffer Overflow Probabilities. In *IEEE Infocom*, 2002.
- [8] J. Choe and N. Shroff. A Central Limit Theorem based Approach for Analyzing Queue Behavior in High Speed Networks. *IEEE/ACM Transactions on Networking*, 6(5):659–671, Oct. 1998.
- [9] M. E. Crovella and A. Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. *IEEE/ACM Transactions on Networking*, 5(6), Dec. 1997.
- [10] T. Dinh, B. Sonkoly, and S. Molnár. Fractal Analysis and Modeling of VoIP Traffic. In *International Telecommunication Network Strategy and Planning Symposium (Networks)*, pages 123 – 130, Vienna, Austria, June 2004.
- [11] C. Fraleigh, F. Tobabi, and C. Diot. Provisioning IP Backbone Networks to Support Latency Sensitive Traffic. In *IEEE Infocom*, Apr. 2003.

- [12] M. Grossglauser and J.-C. Bolot. On the Relevance of Long-Range Dependence in Network Traffic. *IEEE/ACM Transactions on Networking*, 7(5):629–640, Oct. 1999.
- [13] M. Grossglauser and D. N. C. Tse. A Framework for Robust Measurement-Based Admission Control. *IEEE/ACM Transactions on Networking*, 7(3):293–309, 1999.
- [14] O. Heckmann and J. Schmitt. Best-Effort versus Reservations Revisited. In *13th IEEE International Workshop on Quality of Service (IWQoS)*, Passau, Germany, June 2005.
- [15] S. Jamin, P. Danzig, S. J. Shenker, and L. Zhang. Measurement-Based Admission Control Algorithms for Controlled-Load Services Packet Networks. In *ACM SIGCOMM*, 1995.
- [16] F. P. Kelly. *Stochastic Networks: Theory and Applications*, volume 4, chapter Notes on Effective Bandwidths, pages 141 – 168. Oxford University Press, 1996.
- [17] E. Knightly and N. Shroff. Admission Control for Statistical QoS: Theory and Practice. *IEEE Network Magazine*, 13(2):20 – 29, 1999.
- [18] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Transactions on Networking*, 2(2), Feb. 1994.
- [19] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, and C.-N. Chuah. Characterization of Failures in an IP Backbone. In *IEEE Infocom*, Hongkong, Mar. 2004.
- [20] R. Martin, M. Menth, and J. Charzinski. Comparison of Border-to-Border Budget Based Network Admission Control and Capacity Overprovisioning. In *4th IFIP-TC6 Networking Conference (Networking)*, pages 1056 – 1068, Waterloo, Canada, May 2005.
- [21] M. Menth. *Efficient Admission Control and Routing in Resilient Communication Networks*. PhD thesis, University of Würzburg, Faculty of Computer Science, Am Hubland, July 2004.
- [22] M. Menth, J. Milbrandt, and S. Oechsner. Experience-Based Admission Control (EBAC). In *9th IEEE Symposium on Computers and Communications (ISCC)*, pages 903 – 910, Alexandria, Egypt, June 2004.
- [23] K. Murakami and H. S. Kim. Optimal Capacity and Flow Assignment for Self-Healing ATM Networks Based on Line and End-to-End Restoration. *IEEE/ACM Transactions on Networking*, 6(2):207–221, Apr. 1998.
- [24] K. Nichols, V. Jacobson, and L. Zhang. RFC2638: A Two-Bit Differentiated Services Architecture for the Internet, July 1999.
- [25] P. Pan, G. Swallow, and A. Atlas. RFC4090: Fast Reroute Extensions to RSVP-TE for LSP Tunnels, May 2005.
- [26] K. Papagiannaki, N. Taft, Z.-L. Zhang, and C. Diot. Long-Term Forecasting of Internet Backbone Traffic: Observations and Initial Models. In *IEEE Infocom*, San Francisco, CA, Apr. 2003.
- [27] V. Paxson and S. Floyd. Wide-Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.
- [28] A. Pras, R. van de Meent, and M. Mandjes. QoS in Hybrid Networks – An Operator’s Perspective. In *13th IEEE International Workshop on Quality of Service (IWQoS)*, Passau, Germany, June 2005.
- [29] J. Roberts, U. Mocci, and J. Virtamo. *Broadband Network Teletraffic - Final Report of Action COST 242*. Springer, Berlin, Heidelberg, 1996.
- [30] J. W. Roberts. Traffic Theory and the Internet. *IEEE Communications Magazine*, 1(3):94 – 99, Jan. 2001.
- [31] K. W. Ross and D. H. K. Tsang. The Stochastic Knapsack Problem. *IEEE/ACM Transactions on Networking*, 37(7):740–747, 1989.
- [32] T. Schwabe and C. G. Gruber. Traffic Variations Caused by Inter-domain Re-routing. In *International Workshop on the Design of Reliable Communication Networks (DRCN)*, Ischia Island, Italy, Oct. 2005.
- [33] S. Shenker. Fundamental Design Issues for the Future Internet. *IEEE Journal on Selected Areas in Communications*, 13(7):1176–1188, Sept. 1995.
- [34] I. Stoica and H. Zhang. Providing Guaranteed Services without per Flow Management. *ACM SIGCOMM Computer Communications Review*, 29(4), Oct. 1999.
- [35] R. Szábó, T. Henk, V. Rexhepi, and G. Karagiannis. Resource Management in Differentiated Services (RMD) IP Networks. In *International Conference on Emerging Telecommunications Technologies and Applications (ICETA 2001)*, Kosice, Slovak Republic, Oct. 2001.
- [36] H. Tuan Tran and T. Ziegler. Adaptive Bandwidth Provisioning with Explicit Respect to QoS Requirements. In *International Workshop on Quality of Future Internet Services (QoFIS)*, Stockholm, Sweden, Oct. 2003.
- [37] R. van de Meent and M. Mandjes. Evaluation of ‘User-Oriented’ and ‘Black Box’ Traffic Models for Link Provisioning. In *1st Conference on Next Generation Internet Design and Engineering (NGI)*, Rome, Italy, Apr. 2005.
- [38] R. van de Meent, A. Pras, M. Mandjes, H. van den Berg, and L. Nieuwenhuis. Traffic Measurement for Link Dimensioning: A Case Study. In *14th IFIP/IEEE Workshop on Distributed Systems: Operations and Management (DSOM)*, pages 106 – 117, Oct. 2003.
- [39] H. van den Berg, M. Mandjes, R. van de Meent, A. Pras, F. Roijers, and P. Venemans. QoS-Aware Bandwidth Provisioning for IP Network Links. *Computer Networks*, 50(5):631 – 647, Apr. 2006.
- [40] G. van Hoey, D. de Vleeschauwer, B. Steyaert, V. Ingelbrecht, and H. Brunel. Benefit of Admission Control in Aggregation Network Dimensioning for Video Services. In *3rd IFIP-TC6 Networking Conference (Networking)*, pages 357 – 368, Athens, Greece, May 2004.
- [41] Z.-L. Zhang, Z. Duan, Y. T. Hou, and L. Gao. Decoupling QoS Control from Core Routers: A Novel Bandwidth Broker Architecture for Scalable Support of Guaranteed Services. In *ACM SIGCOMM*, pages 71–83, 2000.
- [42] Z.-L. Zhang, V. J. Riberio, S. Moon, and C. Diot. Small-Time Scaling Behaviors of Internet Backbone Traffic: An Empirical Study. In *IEEE Infocom*, San Francisco, CA, Apr. 2003.