| On Burst Assembly in Optical Burst Switching Networks - A Performance	Về Vấn Đề Lập Burst trong Các Mạng Chuyển Mạch Chùm Quang-Đánh giá hiệu

Evaluation of Just-Enough-Time Check 2

On the way towards an architecture for a new QoS-supporting and scalable Internet, the IP- over-photonics approach seems to be very promising. One possible solution in this domain is optical burst switching (OBS), a concept combining advantages of optical circuit and packet switching. After an introduction to OBS as well as the reservation mechanism Just-EnoughTime (JET) we present an approximative analysis of the burst loss probability in an OBS node for an arbitrary number of service classes. Based on analytical and simulation results, we show the impact of traffic characteristics on service differentiation in a single node. Finally, we investigate service differentiation for various parameters in an OBS network scenario.

1.   INTRODUCTION

The current Internet is suffering from its own success. As the number of users on the Internet and the variety of applications transported are growing steadily at high rate, the available bandwidth as well as the best-effort paradigm are facing limits. Ubiquitious and frequent con-gestion situations restrict the use of new time-critical applications like IP telephony, video con-ferencing or online games. Thus, there is not only an increasing demand for bandwidth but also some sort of scalable quality of service (QoS) support. One evolution trend is towards the transport of IP directly over the photonic layer (IP-over-photonic), only with a thin

Trong quá trình hướng đến một kiến trúc Internet khả mở rộng và có cơ chế hỗ trợ QoS mới, phương pháp tiếp cận IP trên photonics có vẻ rất hứa hẹn. Một trong những giải pháp khả dĩ trong lĩnh vực này là chuyển mạch chùm quang (OBS), một phương pháp kết hợp được những ưu điểm của chuyển mạch kênh quang và chuyển mạch gói tin. Sau khi giới thiệu OBS cũng như cơ chế dành trước tài nguyên Just-EnoughTime (JET), chúng tôi sẽ trình bày phân tích gần đúng về xác suất tổn hao gói tin trong một nút OBS ứng với một số lượng lớp dịch vụ tùy ý. Dựa trên kết quả phân tích và mô phỏng, chúng tôi làm rõ tác động của các đặc trưng lưu lượng đến sự khác biệt dịch vụ trong một nút đơn. Cuối cùng, chúng tôi khảo sát sự khác biệt dịch vụ đối với các tham số khác nhau trong trường hợp mạng OBS.

adaptation layer in between [4, 8]. The major advantage of this approach is to reduce overhead caused by overlaid functionality. Furthermore, the success of IP-over-everything is continued while the optical layer provides sufficient bandwidth. Now, the big challenge is to make the optical layer - which currently usually employs static, circuit switched transmission pipes - more dynamic [14]. On this way, two major problems of photonics have to be considered: there is no optical bit processing at high speed and there is no flexible optical buffering beyond fiber delay lines. Therefore, an architecture for the future Internet cannot apply QoS mechanisms ported from electrical networks but should take advantage of photonic network properties.

Three main approaches for a more dynamic photonic layer with QoS support are optical label switching (OLS, including MPLS [13], MP1S [3, 9] and GMPLS [2]), OBS [12, 16, 17] and optical packet switching (OPS) [15]. While OLS provides bandwidth at granularity of a wavelength OPS can offer an almost arbitrary fine granularity, comparable to currently applied electrical packet switching. OBS, which is described in the following, lies between them.

The remainder of this paper is organized as follows: section 2. introduces the functionality and design issues of OBS and shortly resumes the reservation mechanism JET that allows serv-ice differentiation. In section 3. an approximative analysis for the burst loss probability for an arbitrary number of classes and arbitrary QoS offsets is presented. In section 4. we evaluate the performance of different scenarios by analysis and simulation. The focus lies on burst charac-teristics resulting from an assembly process at the edge of the optical network. Furthermore, we discuss service differentiation for various parameters in an OBS network scenario.

2. OPTICAL BURST SWITCHING (OBS)

2.1. Definition and motivation of OBS

Recently, OBS was proposed as a new switching paradigm for optical networks requiring less complex technology than packet switching. OBS is based on some concepts developed several years ago for electronic burst switching networks. At that time, burst switching essentially was an extension of fast packet switching with packets of variable and arbitrary length employing decentralized shared buffer switches [1]. The main characteristics of OBS are the hybrid approach of out of band signalling and electronic processing of header information while data stays in the optical domain all the time, one-pass reservation, variable length bursts, and no mandatory need for buffers.

In principle, burst transmission works as follows (Fig. 1): arriving IP packets are assembled to bursts at the edge of the OBS network. Hereby, the assembly strategy is a key design issue on which we elaborate in section 4.. Transmission and switching resources for each burst are reserved according to the one-pass reservation scheme, i.e. data is sent shortly after the reser-vation request without receiving an acknowledgement of successful reservation. On the one hand, bursts may be released into the network although there are not enough resources available and therefore be lost, on the other hand, this yields extremely low latency as propagation delay usually dominates transmission time in wide area networks. The reservation request (control packet) is sent on a dedicated wavelength some offset time prior to the transmission of the data burst - we classified this as separate-control delayed-transmission (SCDT) in [5].

This basic offset has to be large enough to electronically process the control packet and set up the switching matrix for the data burst in all nodes. When a data burst arrives in a node the switching matrix has been already set up, i.e. the burst is kept in the optical domain.

arrival of control packet

Figure 2. Reservation scenario for bursts of different classes

2.2. The reservation mechanism Just-Enough-Time (JET)

Concerning reservation of wavelengths for burst transmission, different protocols are pro-posed that can be classified as SCDT. In [5] we give a detailed overview, classification and per-formance comparison of the most important proposals. A reserve-a-fixed duration (RFD) scheme reserves all resources exactly for the transmission time of the burst. JET is a RFD scheme proposed by Qiao and Yoo in [12]. Here, predetermined start and end times of each burst are considered for reservation. First, this allows to efficiently use resources, second, it allows for service differentiation by an additional (QoS) offset for higher priority classes. A larger offset permits a higher priority class of bursts to reserve resources in advance of a lower priority class with a shorter offset. However, as larger offsets cause additional fixed delay this offset time has to be carefully chosen. Fig. 2 illustrates a scenario with three wavelengths where a high and low priority burst arrive at the same time. It can be seen that the low priority burst cannot be served as all wavelengths are already occupied during its transmission time whereas the high priority burst is able to find a wavelength due to its much larger offset.

2.3. Key design parameters of a JET-OBS network

OBS and the just introduced reservation protocol JET offer a variety of parameters

[11]. Some of them can be chosen almost arbitrarily whereas others directly depend on technology. Among the arbitrary parameters are number of classes, burst length distribution (including mean value) and QoS offset to separate classes. Main technological parameters are number of wavelengths and basic offset to compensate processing and switching times. Section 4. discusses the impact of these parameters on performance.

3. PERFORMANCE ANALYSIS

In this section, we present an analysis of the burst loss probabilities of a JET-OBS node, that distinguishes multiple classes of equal mean burst length, for arbitrary offsets. The loss proba-bility is calculated for an WDM output link assuming full wavelength conversion capability. In section 3.1. we start with two classes and extend the analysis to multiple classes in section 3.2..

Unlike the single class case where all bursts have the same fixed basic offset §b to compen-sate switching and processing times we follow - as mentioned in section 2.2. - [18] to introduce additional offsets for all but the least priority class, called QoS offset 5goS, that provide service class differentiation. For the following analysis, we assume that class i has priority over class j if i< j for positive i, j, i.e. our highest priority class has index 0.

One motivation why some network should support only two classes - e.g. stream and elastic- is the debate in the Internet community and recent results indicating that this QoS support might be

sufficient [6]. However, even in a network scenario with only two service classes the reduction of the basic offset in each node to account for experienced processing delay effec-tively leads to the multi-class case. Then, bursts are additionally distinguished based on the number of links still to traverse to their destination.

If the basic offset and all QoS offsets are constant the degree of isolation between two arbi-trary classes solely depends on their effective offset difference, i.e. the constant basic offset has no impact on isolation. This stems from the fact that a constant basic offset db for all classes can be interpreted as a constant shift in time of the reservation process and thus neither arrival nor reservation events are reordered in time. This result has also been proven by simulation for various arrival and service time distributions and offsets. Hence, we assume db = 0 without loss of generality and introduce the effective offset difference D, j between class i and j as

(1)

3.1. Single node with two classes

3.1.1. Basic formulae

Under the assumption that control packets (and thus data bursts) arrive in a Poisson stream we can use Erlang's well-known B formula for the loss probability of a M/G/n loss system

(2)

for an overall offered load A and bundle size n. In [18] it has been shown by

simulation that the conservation law is satisfied for an OBS system with equal mean burst length. If this con-servation law holds, the overall burst loss probability Ploss au is not dependent e.g. on the number of classes. Thus, Ploss au on the considered output link in a two-class OBS node with total offered load A0 + A1 can be obtained independent of service differentiation as

(3)

In order to calculate the burst loss probability of the high priority class Ploss 0 , not only the offered load A0 of the high priority class has to be considered but also a fraction of the carried traffic of the low priority class. This low priority traffic (A0 1) represents bursts which started transmission prior to the arrival of the high priority control packet and are still being served when the high priority burst starts, i.e. A0 1 after the high priority QoS offset began. This additional traffic stems from the fact that high priority traffic is not totally isolated from low priority traffic. Thus, PLoss, 0 is approximated by

(4)

The burst loss probability of the low priority class Ploss 1 can be obtained solving

(5)

with arrival rates l0 and l1 for this output link, respectively. This averaging weights burst loss probabilities with respect to their occurrence. For the carried traffic ^1(A0 1) we have

(6)

where A1 • (1 - Ploss 1) is the carried traffic of the low priority class at the time

when the high priority control packet arrives. $1 - F_f(\Delta_0^1)$ is the complementary distribution function of the forward recurrence time of the burst transmission time at time Δ_0^1. It describes the probability that a low priority burst that has already started transmission prior to some random observation time t has not finished transmission within the period $[t, t + \Delta_0^1]$. In our case, this observation time corresponds to the arrival time of a high priority control packet. Finally, (6) is an approximation because in reality, longer bursts are discarded with a higher probability [5].

3.1.2.Iterative solution

According to (4), (5) and (6), there is a mutual dependency between P_{loss}^0 and P_{loss}^1. Therefore, we suggest an iterative solution for above formulae. We initialize the iteration with estimates for loss probabilities of high and low priority classes, $\overline{P}_L^{ss}{}^0$ and $P]0^1$. These zero order estimates are given in (7) and can be derived from (3) - (5) by decoupling the high priority class from the low priority class which is equivalent to neglecting (Δ_0^1).

(7)

Similar formulae are also published by Qiao and Yoo [18] and yield lower boundaries for our analysis if the QoS offset is very large (Fig. 5, see below).

The distribution function of forward recurrence time of burst transmission time is given by

(8)

where $\hat{\mu}$ and $\hat{F}(u)$ represent mean and distribution of the burst transmission time, respectively. Finally, the amount of carried low priority traffic is determined by (6) using (7) and (8)

$$(9)$$

and can be inserted in (4) yielding a first order result for the loss probability of the high priority class $P_{Loss,0}$. By application of the conservation law (5) and the just derived result for $P_{Loss,0}$ a first order result for the low priority class $P_{Loss,1}$ is obtained. Iteration until some precision criterion is satisfied leads to $P_{Loss,0}$ and $P_{Loss,1}$.

3.2. Single node with arbitrary number of classes

3.2.1. Basic formulae

The burst loss probabilities for k service classes with different QoS offsets is obtained by heuristically generalizing basic formulae (3) - (6) to an arbitrary number k of classes. This is performed by considering all interference from a class m of lower priority on a class i of higher priority ($0 < i < m < k- 1$). $P_{Loss,all}$ again follows Erlang's loss formula as given in (2). $P_{Loss,0}$ is calculated by taking into account its own offered load A_0 and the interfering carried traffic components $Y_m(A_0)$ originating from lower priority class m

$$(10)$$

In the multi-class case, a conservation law corresponding to (5) can be formulated for every set of classes $S_j - \{0, j\}$ with $0 < j < k- 1$

$$(11)$$

where $P_{Loss,S}$ is the total loss probability of all classes in S_J . Each class i in S_J

experiences additional inteerfering traffic Ym(Ai m) from each class m not belonging to SJ

Figure 3. Network scenario with reference path

(12)

These interference components are weighted by the arrival rate of class i within S, - representing relative occurrence of class i bursts in S, - and summed up over all i and m for given j

Consequently, (10) and the set of k - 1 equations in (11) completely describe approximations of burst loss probabilities for all k classes.

3.2.2.Iterative solution

Again, we suggest the iterative solution of (10) - (13). Starting with (10) for the highest pri-ority class, we repeatedly solve (11) for Plossj with increasing class indices J. We calculate initial values for PL0)ss 0 from (10) and for all other PLossj from set of equations (11) assuming no interference, i.e.

(14)

These zero order estimates have been described in [18]. They yield lower boundaries in case of perfect isolation with a'. . + 1 ® ¥ , i.e. no interference of classes. By evaluating (12) for zero order estimates and inserting results in (10) and (11) first order results for all Plossj can be cal-culated. Iteration until some precision criterion is satisfied leads to all burst loss probabilities.

3.3. Application to multiple nodes

In order to apply the above presented

theory to an OBS network, we suggest to apply the well-known stream analysis which is based on decomposition and assumption of independence (see section 4.2.2.). Fig. 3 shows a reference path through an open queueing network from a source node to a destination node traversing core nodes 1 to N. We start solving the burst loss probability Ploss i(v) for class i at node v on the respective output link with all aggregated arrival rates l, v for all classes J at node v and above presented formulae. By considering the reference path, the arrival rate l f . of class i reduces to l f . • (1 - PLossi(1)) after node 1, href i • (1 - Ploss i(1)) • (1 - Ploss i(2)) after node 2 etc. Hence, after node N, we have

And the end-to-end burst loss probability Peoss ref j for class i on the reference path as

Figure 4. Impact of arrival process on burst loss probabilities (A0 j/hx = 1

Figure 5. Analytical and simulation results for high priority burst loss probability

4. EVALUATION OF SERVICE DIFFERENTIATION CAPABILITY

In section 4.1. we regard a single isolated node, while the focus in section 4.2. is on multiple nodes in a network scenario. For the following evaluations, we assume the number of wave-lengths to be 8 in a two-class OBS system with a relative high priority traffic share of 30% at a total load of 0.6. Restriction to 8 wavelengths allows us to perform simulations with sufficient accuracy in acceptable time. Nevertheless, as simulation and analysis have proven to match well we can obtain

results by analysis for a higher number of wavelengths. We have shown in [7], that principle effects and shape of curves remain unchanged while the order of magnitude of characteristic values changes.

4.1. Impact of traffic characteristics

In this section, we investigate system performance for different burst characteristics in order to specify requirements and trade-offs for assembly strategies. First, we address the impact of burst interarrival time distributions, then we focus on burst lengths. Pre-transmission delay faced by a high priority, potentially real-time, burst comprises the time until a burst is assembled and a reservation is initiated^ as well as the offset. While the first component is proportional to the actual high priority burst length the latter grows with mean low priority burst length. Thus, assembly strategies have to find suitable burst lengths.

4.1.1. Interarrival time distribution

As the assumption that the burst interarrival time has Markovian property seems to be very restrictive, we carried out simulations varying the burst interarrival time distribution of both classes. In Fig. 4, burst loss probabilities of a high and a low priority class for different uncor-related interarrival time distributions^ and negative-exponentially distributed burst lengths (ho = hj) are depicted against the load. It can be seen that changes in the arrival process have only small impact on the burst loss

probabilities of both classes. Thus, the model of a Poisson arrival process yields reasonable results even for very different interarrival time distributions.

t This pre-transmission waiting time could be reduced by intelligent algorithms for initiating reservation control packets ahead of time. Imperfect prediction regarding burst length, however, leads to overhead due to waisted bandwidth.

t The hyperexponential distribution satisfies the symmetry condition p hj = (1 - p) ' h2 where p is the branch probability and hi and h2 are the mean values of the respective phases.

ratio of mean burst lengths of classes 0 and 1

Figure 6. Impact of mean burst length on mean burst loss probability

4.1.2. Burst length distribution

In this section we assume mean burst transmission times of high and low priority bursts to be the same. Fig. 5 shows Peoss 0 against the QoS offset normalised by h1 for different low priority burst length distributions. An upper boundary for the case of no isolation as well as a lower boundary for perfect isolation (see section 3.1.2.) are included. It can be seen that our presented analysis matches the simulated curves quite well for all distributions. The strong impact of the forward recurrence time of the low priority burst length as indicated by (4), (6) can be observed. A hyperexponential distribution for the low priority burst length with high coefficient of variance (CoV) leads to a significant

increase of Peoss 0 and a very slow approach of the lower boundary even for large QoS offsets. Nevertheless, in contrast to intuition, CoV is not the decisive factor as can be observed for the Pareto distribution which has hardly any impact compared to negative-exponentially distributed low priority bursts. In case of small offsets Pareto distributed burst lengths even yield better performance. Thus, the assembly strategy has to carefully shape low priority bursts in order to efficiently operate the system. As we showed in [7], the principal shape of curves shown in Fig. 5 remains unchanged for an increasing number of wavelengths. Only the order of magnitude of losses changes drastically, e.g. for 64 wavelengths the lower boundary reduces to about 10 . In all following evaluations, we only show results for negative-exponentially distributed burst lengths.

4.1.3. Mean burst lengths

In order to reduce processing overhead and increase efficiency for large volume bulk traffic longer low priority bursts might be advantageous. However, in order to maintain a certain degree of isolation, larger low priority bursts result in a larger QoS offset and consequently a longer pre-transmission delay for the high priority class. With respect to this trade-off, we evaluate the performance of an OBS node depending on the ratio of the mean burst lengths h0 1 = h0/h1. In order to keep the offered load Ai = 1i • hi unchanged within each class we adapt the arrival rates. Fig. 6 shows Peoss an

against h0 1. In this graph curves are drawn for several offsets. As expected, Peoss an is unchanged for varying ^ 1 if no offset distinguishes the classes. But even for very small offsets - e.g. introduced by basic offsets, see also section 4.2.- Peoss au changes significantly with h0 1. For shorter high priority bursts Peoss au decreases while it increases for longer high priority bursts. Thus, a decreased PLoss all can be achieved by operating the system with bursts satisfying h0 1 < 0.7. This scenario contradicts

ratio of mean burst lengths of classes 0 and 1 Figure 8. Impact of mean burst length on high priority burst loss probability burst transm. time / mean low priority burst transm. time

Figure 9. Low priority burst loss probability conditioned on burst length the conservation law and therefore is not covered by our analysis. However, as indicated in Fig. 8, Ploss o hardly changes over hQ j and is thus still reasonably approximated by

In order to get a deeper inside into this effect, the burst loss probabilities of both classes are observed separately by simulations with the same parameters as in Fig. 6. From Fig. 7 it can be seen that PLOSS 1 significantly increases for decreasing ho j. This effect is caused by the reser-vation mechanism itself, as low priority bursts in most cases fill gaps left over by high priority bursts. Due to the higher number of arriving high priority bursts per time interval, the link is

fragmented and the length of gaps left for low priority bursts is reduced. This explanation is confirmed by Fig. 9 where PLOSS 1 is depicted conditioned on the low priority burst length for different values of ho j. The QoS offset is chosen corresponding to DO j/hj = 1. It can be seen that the burst loss probability increase is larger for lower ho j. If the burst transmission time is longer than the offset duration, a boundary value is reached, which we showed in [5]. This boundary value increases for decreasing ho j. Again, very short bursts are not affected as they fit into small gaps left over.

Resuming the above discussion, Fig. 8 indicates that Ploss o slightly decreases for shorter high priority bursts. Together with the description of Ploss o in (4) and (6) and the increase of Ploss 1, the decrease of Ploss o can be explained: High priority traffic experiences reduced low priority interference due to higher low priority losses. Considering the significant changes of the arrival rates over ho j in (5) as well as the behaviour of Ploss o and Ploss j, the dependence of Ploss au on ho j depicted in Fig. 6 can now be explained.

Summarizing, on the one hand, it is desirable to have a small ho j because it fits the idea of short high priority, potentially real-time bursts and long bulk traffic low priority bursts, and it results in a reduced PLOSS all. On the other hand, if ho j is small, PLOSS j increases significantly for longer low priority bursts. This is undesirable, especially as

from the signalling and processing point of view, it is much more efficient to transmit long low priority bursts.

4.2. JET in an OBS network scenario

In the following, we discuss the burst loss probabilities in a simple network scenario where every destination can be reached with either one or two hops. This is reasonable for a future national core network in a country like Germany [10]. In section 4.2.1., we look at effects in a single node in a network scenario while we look at network wide effects in section 4.2.2..

QoS offset / mean transmission time Figure 10. Comparison of analytical and simulation results for two-class network scenario

4.2.1. Multiple effective classes due to basic offset adaptation

In a network scenario, bursts with a different number of remaining hops to their destination have different basic offsets as the offsets are decreased in every OBS node traversed. The resulting differentiation based on QoS as well as basic offset can be described by an increased number of effective classes. Approximations of burst loss probabilities for the effective classes can be calculated with the multi-class analysis presented in section 3.2.. For two service classes in a two hop network, i.e. bursts have either one or two more nodes to traverse (as in Fig. 12), four effective classes have to be considered.

In order to get an idea how basic offset δ_b, QoS offset δ_{qos}, and mean burst length should be chosen, we introduce a basic offset ratio as $r_h = \delta_b/\delta_{qos}$. While δ_b is

determined by the speed of processing and switching, 5qos can be chosen rather independently always keeping in mind its influence on loss probability and delay. Original traffic flows and classes are mapped to effective classes according to Table 1. In Fig. 10 and Fig. 11 burst loss probabilities are depicted for different values of rb against &QOS/h1. In Fig. 10, we compare analytical and simulation results for rb = 0.1. It can be seen, that the shapes of respective curves match rather well and that the following principle effects are described by the analysis. From Fig. 11, it can be observed that the curves diverge for both increasing 5qos and increasing rb. However, an increased rb significantly splits up both, the high priority class and the low priority class, which is very undesirable as bursts which already occupy resources are discriminated. For instance, high priority bursts of the two hop flow at their last hop (effective class 2), which already occupy resources on their first hop link, have a higher loss probability than any high priority burst at its first hop (effective class 0). Thus rB < 0.1 must hold in order to keep the difference in loss probabilities to roughly less than one order of magnitude for QoS offsets

2 hop flow at first hop(eff. class 0 and 1)

2 hop flow at last hop (eff.class 2 and 3)

1 hop flow atlast hop (eff. class 2 and 3)

Figure 12. Traffic flows and effective classes at the evaluated node

Table 1. a. Flows and classes b. Effective

classes

s Ratio of "through" and "local" traffic load Figure 13. Burst loss probability in a tandem model with varied „through traffic"

5qos < 3 • hi and to allow a reasonable operation in a multi-hop environment. For rB > 0.1 or very large offset values, this spreading in more classes has to be avoided by placing a fiber delay line of length §b in front of each JET-OBS node. This fiber delay line compensates processing and switching times and makes a basic offset unnecessary.

4.2.2. Generalization of single-node results to networks

In this section, we study the assumption that congestion in an OBS-node is independent of the origin of traffic streams as long as they are mixed to a certain degree. If a stream of bursts traverses a sequence of nodes without injection of any other bursts there will be no blocking but in the first node. However, if traffic leaving a node is split up among several nodes and input traffic into a node comprises traffic from several preceding nodes, blocking is almost equal for all streams. In Fig. 13 we varied the ratio of traffic which has already undergone a reservation process in a preceding node (through traffic, e.g. solid line at second node in Fig. 12) and traffic which has not (local traffic, e.g. dashed lines at second node in Fig. 12) and plotted the ratio of loss probabilities of through and local traffic. It is shown that for a smaller traffic share of an individual traffic stream, the loss ratio increases and

approaches 1.

In a meshed core network we assume node degrees of at least four (splitting ratio < 0.33 in Fig. 13) allowing the approximation of independent loss probabilities. Due to this justification we can apply the results for the single-node evaluation also to OBS networks as proposed in section 3.3.. The end-to-end loss probability can be estimated by the solution given in (16).

5. CONCLUSIONS AND OUTLOOK

An overview of optical burst switching (OBS) and the reservation mechanism Just-EnoughTime (JET) is provided. We presented an approximative analysis to calculate the burst loss probability for an arbitrary number of classes and arbitrary offset values in an OBS node. By this analysis as well as a simulation tool, we evaluated the performance of an OBS node in different scenarios. Thereby, we found out that this reservation protocol is strongly dependent on burst characteristics resulting from burst assembly at the edge of the optical network. The differentiation of classes not only depends on the burst length distribution function, but also on the ratio of the mean burst lengths of the classes. Nevertheless, a good degree of QoS can be achieved applying JET if the burst assembly strategy produces proper burst characteristics. In a network scenario, the ratio of basic offset compensating switching and processing delay and

QoS offset differentiating classes has a strong impact on intra-class differentiation and therefore has to be kept well below 0.1. Our presented multi-class analysis covers this behaviour by considering an increased number of effective classes.

Further work should include the design, implementation and evaluation of assembly strategies based on dependencies on burst characteristics presented here. Furthermore, optimization of the reservation mechanism that improve the transport of very long low priority bursts are desirable. Finally, the impact of partial wavelength conversion capabilities has to be studied.