Optical Switching and Networking 24 (2017) 47–56

Contents lists available at ScienceDirect

Optical Switching and Networking

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

Providing efficient support for real-time guarantees in a fibre-optic AWG-based network for embedded systems

Kristina Kunert*, Magnus Jonsson, Annette Böhm, Tomas Nordström

Halmstad University, Box 823, SE-301 18 Halmstad, Sweden

ARTICLE INFO

2000 MSC:
00-01
68M10
68M20

Keywords:
Real-time communication
Embedded system
Worst-case analysis
Multiwavelength network
Arrayed waveguide grating

ABSTRACT

High-performance embedded systems running real-time applications demand communication solutions providing high data rates and low error probabilities, properties inherent to optical solutions. However, providing timing guarantees for deadline bound applications in this context is far from basic due to the parallelism inherent in multiwavelength networks and often bound to include a large amount of pessimism. Assuming deterministic medium access, an admission control algorithm using a schedulability analysis can ensure deadline guarantees for real-time communication. The traffic dependency analysis presented in this paper is specifically targeting a multichannel context, taking into consideration the possibility of concurrent transmissions in these types of networks. Combining our analysis with a feasibility analysis in admission control, the amount of guaranteed hard real-time traffic could be shown to increase by a factor 7 in a network designed for a radar signal processing case. Using this combination of analysis methods will render possible an increased amount of hard real-time traffic over a given multichannel network, leading to a more efficient bandwidth utilization by deadline dependent applications without having to redesign the network or the medium access method.

1. Introduction

As the usage of embedded systems propagates into new application areas, requirements on embedded communication networks continue to increase. New network architectures are proposed and new components are employed to meet demands like high performance, speed and reliability at a reasonable price. Real-time services, i.e., being able to guarantee that the data over a communication network reach their destination within a predefined deadline, play a vital role in high-performance networks applied to, e.g., cluster computing [1], radar signal processing [2] and in-vehicle networks [3].

Due to their advantageous properties, optical interconnects and devices are commonly used in these modern communication systems. High bandwidth, low loss, and immunity to electromagnetic interference and radio frequency interference are some of the reasons that give optical architectures an advantage over those merely built upon electronics. While the usage of optics is rather common in long distance communication systems, the research on optical solutions within embedded systems is more limited. Thanks to recent advances in optical components, wavelength-sensitive components, as, e.g., Arrayed Waveguide Gratings (AWGs), nowadays present a viable choice due to their reported good qualities, as, e.g., high concurrency, high reliability, low loss, and non-blocking connectivity. AWG-based networks guaranteeing real-time support, as needed for the aforementioned application areas, is an area that is rather unexplored by researchers so far.

Providing definitive, hard real-time guarantees, i.e., a guarantee that ensures that every deadline is met, is not elementary in a communication context. The key to success is the usage of deterministic medium access control (MAC) and an admission control system not allowing more traffic onto the network as can be served before any given deadline. Deterministic medium access usually is achieved either by prescheduling access or having a control node administering access in run time, e.g., by the usage of a token. Mostly those protocols are based on time division multiple access (TDMA). While the MAC protocol is responsible for decision on which packet is to be sent when, the admission control has to decide which traffic is allowed on the network in the first place. This means it provides the information to the MAC protocol about what packets are to be sent. In order to be able to guarantee any deadlines, a worst case situation (in terms of delay) has to be analysed by the admission control and, depending on the outcome, decisions on which traffic to admit can be made.

The goal of our work is to improve the real-time services optical multiwavelength networks for embedded systems can provide. In

* Corresponding author.
E-mail addresses: Kristina.Kunert@hh.se (K. Kunert), Magnus.Jonsson@hh.se (M. Jonsson), Annette.Bohm@hh.se (A. Böhm), Tomas.Nordstrom@hh.se (T. Nordström).

http://dx.doi.org/10.1016/j.osn.2016.11.004
Received 22 February 2014; Received in revised form 15 November 2016; Accepted 17 November 2016
Available online 23 November 2016
1573-4277/

earlier work [4] we presented an AWG-based wavelength division multiplexing (WDM) network for implementation in high-performance embedded systems with high demands on bandwidth and timing. A MAC protocol was described that offered support for heterogeneous real-time traffic, using Earliest Deadline First (EDF) scheduling. The AWG's spatial wavelength reuse and a combination of fixed-tuned and tuneable transceivers in the end nodes enabled simultaneous control and data transmissions. This network solution was developed for short range communication systems such as, e.g., System Area Networks (SANs) and similar high-performance embedded systems with possible application areas as cluster computing, distributed large routers and distributed video and imaging applications. Its feasibility was demonstrated by a case study in the area of radar signal processing (RSP).

This previous research provided hard real-time guarantees by means of a deterministic MAC protocol and the usage of a feasibility analysis in the admission process. The analysis has its origin in the area of uniprocessor task scheduling in real-time systems and had previously been mapped onto the communication context in [5]. Generally speaking, the timing and resource usage of a task being run on a processor can be analysed the same way as packets being sent over a network link. However, as multichannel networks allow concurrent transmissions, the uniprocessor feasibility analysis introduced unnecessary pessimism not taking into consideration this parallelism.

The contribution of this paper is a generalization of the feasibility analysis, reducing the pessimism that arises when mapping a method targeting systems with a single resource (as e.g. task executing on a single processor) onto systems with parallel resources (as e.g. packets sent over a multiwavelength network). By analysing traffic interdependencies, and taking them into consideration when applying the feasibility analysis in the admission control, the pessimism can be reduced. Our calculations and simulations show that the improved feasibility analysis results in a higher amount of guaranteed hard real-time traffic admitted to the network.

The rest of the paper is organized as follows. Section 2 provides a short introduction to the aspects of the AWG component relevant for this paper, and gives an insight into related research. Section 3 summarizes the earlier work from [4] that forms the basis to this paper. Section 4 discusses the feasibility analysis and how to improve its basic version through traffic analysis. Section 5 provides simulation results illustrating the degree of improvement, regarding throughput guarantees, possible when using the suggested traffic analysis. Section 6 concludes this paper.

2. Related background

As AWG-based WDM networks have been around for a relatively short number of years, the research on providing real-time communication over them is limited. However, extensive research under a long time has been done in the field of broadcast-and-select networks based on Passive Star Couplers (PSC) [6]. Opposed to the PSC, the AWG is a wavelength-sensitive device, a functionality that makes it possible to reach selected parts of the network and increase its concurrency and throughput. Section 2.1 provides a short introduction to the important properties of the AWG, while Section 2.2 shortly describes MAC protocols for AWG-based single-hop WDM star networks. Section 2.3 presents a small selection of MAC protocols with real-time support for non-AWG-based single-hop WDM star networks in order to provide some background on approaches that offer real-time support in similar networks.

2.1. The AWG component

Optical networks enable communication with low bit error probability and high bandwidth. However, the bandwidth of the optical medium is often not fully exploitable by electronic components working at a limited speed. In order to better utilize the optical bandwidth, a

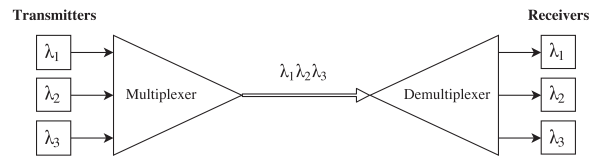

Fig. 1. A simple WDM system.

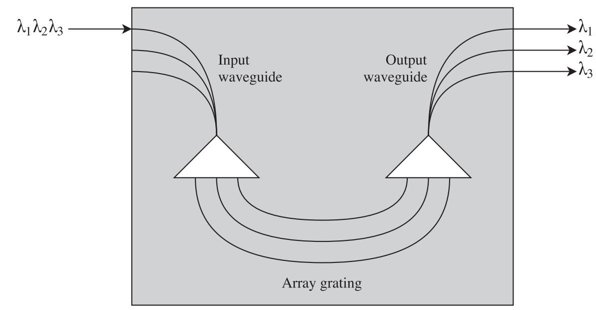

Fig. 2. Illustration of a 3×3 AWG.

technique for implementing concurrency in communication is needed. The most recognized approach is WDM, which increases the transmission capacity by simultaneously transmitting multiple wavelengths through a single optical fibre. In its most basic form this is achieved by a wavelength-sensitive multiplexer combining individual wavelengths onto the fibre, while on the other end a corresponding demultiplexer splits the signals up again and sends them to individual destinations (Fig. 1). In wavelength-routed networks, optical channels are routed to different destinations based on their wavelength. One way of implementing this is by the help of the AWG.

The AWG is an optical component that consists of an arrangement of optical waveguides on a substrate plate (Fig. 2). Its area of usage comprises implementations as wavelength multiplexers and demultiplexers, but also as an integrated building block in WDM equipment [7]. In this paper, the AWG is used as a wavelength router, i.e., a passive routing device for optical communication networks.

An $N \times N$ AWG has N input and N output waveguides (ports). Due to spatial wavelength reuse, $N \times N$ AWGs can simultaneously accept N wavelengths at each of their input ports and, using periodic wavelength routing, route each of these wavelengths to a specific output port (dictated by the physical design of the AWG) without collision. Each output port will in its turn receive N wavelengths, one from each input port. This makes the AWG a strictly non-blocking device and enables the implementation of multiple, concurrent communication channels. Assuming tuneable (or an array of fixed-tuned) transmitters and receivers at the input and output ports, the AWG offers full connectivity on all available wavelengths (Fig. 3). This implies that each node attached to any of the input ports of the component can send on all of the AWG's wavelengths and can directly reach each node connected to any output port. This property makes the AWG very well suited for implementing single-hop networks, where each destination is always only one hop away from each source. In order to make optimal use of the AWG, each directly connected node in the network has to be equipped with transmitters and receivers that support the range of all wavelengths used for communication in the network.

2.2. MAC protocols for AWG-based single-hop WDM star networks

This subsection provides a short summary on a number of MAC protocols presented for AWG networks using WDM to achieve multichannel communication for single-hop star networks. Special consideration is given to the offering of real-time guarantees or other Quality of Service (QoS) supported.

A MAC protocol for a network architecture based on tuneable transmitters and receivers at each node and using time division

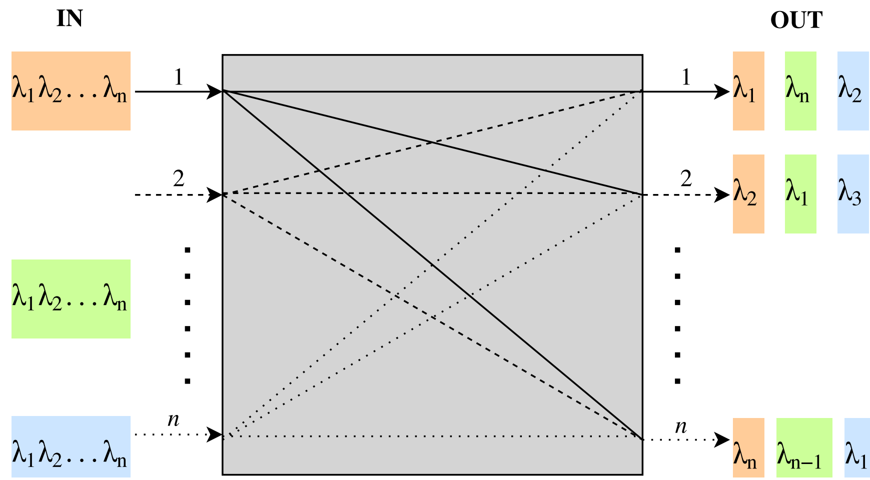

Fig. 3. Wavelength routing pattern of $N \times N$ AWG.

multiple access (TDMA) is proposed in [8]. The protocol guarantees a fair bandwidth allocation with the same amount of bandwidth for each sender-receiver pair, but due to the fixed allocation of capacity to each node, bandwidth utilization might be low while traffic still can suffer from unnecessary delay. This protocol is improved in [9] by introducing priority classes to support traffic with timing constraints, but only the mean queuing delay for each packet is analysed.

A general, collision free MAC protocol for wavelength-routed passive optical networks with tuneable transmitters and fixed receivers is presented in [10], where a minimum bandwidth between each source-destination pair can be guaranteed. However, no deterministic delay bound is given.

By using array components in the end nodes, as proposed in [11], a higher degree of spatial wavelength reuse is allowed for. Only mean packet delay is analysed probabilistically and, as one of the suggested MAC methods is contention-based, no deterministic real-time service can be offered. Similar networks and protocols where fixed medium preallocation limits possible bandwidth utilization for real-time traffic include [12–15].

A reservation-based MAC protocol using an in-band control channel to broadcast control information is introduced in [16]. Data packets are only permitted medium access when the control packet claiming this transmission has been transmitted successfully. However, no evaluation of the protocol is provided, so no decision on its real-time performance can be made.

A frame-based algorithm that provides dynamic bandwidth allocation (DBA) and supports QoS for data, voice and video traffic over a AWG-based Ethernet Passive Optical Network (PON), was designed in [17]. Each frame is divided into dedicated subframes for the three traffic classes of the differentiated services (constant bit rate, variable bit rate, best effort). This ensures a fair division of bandwidth between the different service classes; however, this also increases the possibility of deadline misses when traffic with short deadlines has to wait until the next dedicated frame. Due to the non-deterministic nature of some of the traffic in the study, no deterministic deadline guarantees are targeted, but merely mean packet delay and delay jitter are studied.

More recently, protocols for AWG-based networks have been designed for application within intra-datacentre networks. [18] describe a dynamic wavelength allocation (DWA) technique for an intra-datacentre system area network. Using a multicast-capable AWG-router and multiple tuneable transmitters at each node, the DWA algorithm adapts the number of active transmitters at each node to the current traffic demands. This method shows promise to achieving its goal of providing an energy-efficient allocation technique, but no consideration is taken to possible QoS demands of the traffic. [19] concentrate on investigating the feasibility of an AWG-based load balancer, using wavelength routing for high-performance data centre networking. Targeting guaranteed throughput and delay bounds, and given the underlying all-optical router, load-balancing is necessary to avoid random packet drops due to optical buffer overflow. However, the results indicate decreased bandwidth utilization in order to avoid exceeding buffer space. The usage of First Come First Serve (FCFS) queuing will lead to head-of-line blocking, where packets with long deadlines can block those with shorter deadlines and thereby increase the probability of deadline misses.

None of the aforementioned protocols in this subsection offers any worst-case delay bound analysis, and therefore no real-time guarantees, not even pessimistic ones, can be provided by any of them.

2.3. Real-time MAC protocols for single-hop WDM star networks

In contrast to AWG-based networks, WDM networks based on other optical components have received more attention and a large number of papers have suggested different solutions for traffic with real-time demands. This subsection presents a small selection of them.

In 1995, the first MAC protocol for a single-hop WDM star especially developed for multipriority traffic was presented in [20]. The protocol implements a distributed scheduling algorithm where the packet with the highest priority gets the earliest access to the medium. Reservations are made over a control channel. Unfortunately, the access to the slotted control channel is organized by random access, making deadline guarantees not possible.

In [21–24], preallocation protocols are described which suffer from their inherent disadvantages, as, e.g., potentially low bandwidth utilization. Of those only in [22] a bounded latency for hard real-time traffic is provided, while [24] targets jitter-sensitive applications as, e.g., multimedia streaming and therefore prioritizes constant jitter over guaranteeing hard deadlines.

A hierarchical scheduling framework for on-demand traffic flow scheduling is introduced in [25]. Each node decides on the flow for which to request transmission and, by using control channel communication, a distributed decision is taken on which packets are allowed to transmit. However, the flow scheduling is merely divided into real-time and non-real-time traffic, not taking into account individual flow deadlines. This limitation leads to a higher probability of deadline misses. Additionally, the use of certain random decision-making and an unknown number of possibly necessary packet retransmission makes a deterministic analysis impossible.

Assuming a time-slotted network with support for virtual circuits, [26] develop an algorithm that provides fair network access even to lower priority traffic by assuming that only a certain amount of the network's resources can be allocated to guaranteed traffic.

However, some best-effort traffic might therefore be prioritized over real-time traffic, increasing the risk of missing deadlines. Even [27] designs a protocol taking into consideration a fairness aspect, providing a reservation-based MAC protocol differentiating real-time and non-real-time traffic. Although providing collision-free medium access for both traffic types, FCFS scheduling in the nodes also here introduces a head-of-line effect. Another reservation-based medium access solution providing guaranteed real-time performance is presented in [28].

The mechanism includes an admission control functionality, which, by using a schedulability test, only accepts new traffic into the network if its deadline can be guaranteed and traffic already in the network will not miss its own deadlines due to the increase in traffic. A delay bound analysis is given, but simulations show that only allowing guaranteed traffic onto the network results in poor bandwidth utilization.

In [29,30] they develop a QoS prediction scheme that makes it possible to judge whether traffic with certain QoS demands can be accepted into the network. This approach has certain similarities to our real-time analysis, but while we use our method in connection with admission control to provide deadline guarantees for each packet, the scheme in [29,30] cannot provide any guarantees on the packet level, but merely per traffic class. A different QoS prediction scheme was presented in [31], but merely probabilistic calculations are provided and the accuracy of those predictions is evaluated. Obviously, no timely treatment of hard real-time traffic can be guaranteed this way.

For energy saving purposes, nodes in PONs might enter sleep mode during certain intervals, jeopardizing QoS requirements of delay-sensitive traffic. [32] investigate QoS provisioning in the presence of sleep modes in network nodes that are using the interleaved polling with adaptive cycle time (IPACT) signalling protocol. They provide a queueing analysis of the mean packet delay for the case of each node having access to one WDM channel. No analysis of the multichannel case is provided. Their results include an approximation of the length of the sleeping phase of each node, given a maximum allowed mean packet delay. Due to only studying mean packet delay and not hard deadlines, no guarantees for hard real-time traffic can be achieved with this method.

SpaceFibre [33,3,34,35] is a protocol stack targeting spacecraft applications. A dedicated layer provides QoS by using virtual channels, together with packet prioritization, bandwidth reservation and scheduling/flow control, over a time-slotted medium. Virtual channels can reserve a certain amount of bandwidth and channel precedence (a priority measurement) increases while the channel is not sending any packets. High priority channels always take precedence over lower priority channels. Scheduling can be adapted depending on the demand of the application, rendering possible a completely deterministic mode of operation. However, this deterministic scheduling makes inefficient use of the available bandwidth and scheduling decisions are made independent of traffic properties.

More recently, [36] designed a reservation-based MAC protocol with in-band control signalling. The protocol employs dedicated control slots and packet-by-packet scheduling, which would make it possible to provide real-time services to time-sensitive traffic. However, merely round-robin scheduling independent of packet deadline is implemented and the queueing delay analysis only calculates the mean packet delay. Therefore no real-time guarantees can be given at all.

Other WDM star networks targeting real-time support or providing certain QoS, but without any framework for worst-case delay analysis, include [37–45].

3. The real-time AWG network and MAC protocol

In order to simplify the understanding of our improved feasibility analysis, this chapter of the paper will shortly summarize the relevant facts of the underlying system. This network architecture and MAC protocol, together with a simplistic timing analysis and a case study was introduced earlier in [4]. For a more detailed specification we refer the interested reader to the mentioned publication.

3.1. The AWG network

This chapter introduces shortly the network requirements, the different types of traffic supported by the network, and the network architecture.

3.1.1. Real-time traffic classes

Many embedded systems applications require support for heterogeneous traffic with three general levels of delay sensitivity: hard real-time (HRT), soft real-time (SRT), and non-real-time (NRT). HRT has to meet strict timing requirements. Failing to meet deadlines can have serious consequences for the entire system and a transmission that misses its deadline has lost its value. It is therefore vital to be able to guarantee the success of a transmission in advance (under the assumption of error-free transmission), which is done by the help of the schedule provided by the MAC protocol. SRT also has deadline constraints, but meeting them is not vital to the system. After having transmitted all of HRT traffic, the remaining bandwidth can be used for SRT data. Timeliness is desirable but does not have to be guaranteed. NRT can be described as traffic with an infinite deadline. Those messages can, e.g., be queued until the network capacity allows them to be sent without disturbing any (hard or soft) real-time traffic.

3.1.2. Network architecture

The assumed network architecture consists of one protocol processor (PP), residing with an $N \times N$ AWG in a hub at the centre of a physical, single-hop WDM star topology, and N−1 end nodes (Fig. 4). Single-hop networks provide a direct connection between all source-destination combinations, which makes them easily managed, as well as time and bandwidth efficient. Simultaneous transmissions of control information from all end nodes achieve short control delays in the network. All N nodes, including the PP, are connected to the AWG by two fibres, one for transmission and one for reception. All fibres are assumed to be of equal length, resulting in the same propagation delay for traffic between any source-destination pair.

The network employs a combination of fixed tuned and tuneable transceivers. In each end node a tuneable transceiver covers N−1 wavelengths used for data traffic. One unique wavelength per node is saved for control packet exchange with the PP through the fixed-tuned transceiver. The remaining $N \times 1$ wavelengths at each end node can, at each instance, be used for transmission of data traffic. The AWG routes the packets from source to destination based on the chosen wavelength.

All intelligence about real-time traffic scheduling is embedded in the PP, while the N−1 remaining nodes and the AWG component are completely passive in this respect. The end nodes are merely responsible for scheduling their own output queues. Each of the N−1 end nodes has two different message queues, one for HRT and one for SRT and NRT. The nodes sort their messages in the two queues according to deadline, i.e., the message with the shortest deadline will be placed first in its queue and thereby assigned the highest priority. The EDF scheduling algorithm is chosen due to its optimality under certain circumstances [46].

The PP transmits and receives data via wavelength array compo-

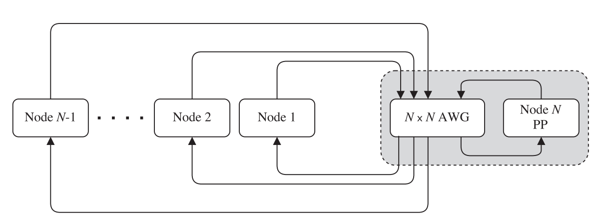

Fig. 4. Network architecture.

nents. The array component on the transmitting side of the PP is attached to a wavelength-sensitive $(N-1)\times 1$ combiner, which multiplexes the N–1 different wavelengths onto one fibre. Correspondingly, a wavelength-sensitive $1\times(N-1)$ splitter separates the incoming wavelengths before they enter the $N\times 1$ array of receivers of the PP. Due to the fixed size of the array components used at the PP, it is not possible to dynamically expand the number of nodes in the network. This lack of scalability is, however, not considered to be a problem in this case, since the target applications are embedded systems with networks that normally do not have to grow during operation.

The fixed-tuned and tuneable transmitter on each of the N–1 remaining nodes are connected to a wavelength-sensitive 2×1 combiner in each node which multiplexes the two wavelengths (one fixed wavelength to the PP and a tuneable one for data transmission to any of the remaining nodes) onto one fibre. The corresponding design on the receiving side of each node is a wavelength-sensitive 1×2 splitter.

3.2. The medium access control protocol

The PP employs a MAC protocol in order to organize the data traffic over the network. The main task of the MAC protocol is to delegate the available wavelengths between the nodes based on control information collected by the PP, and to ensure that the HRT packet with the shortest deadline is sent first.

3.2.1. Control traffic and scheduling

Both control and data traffic are organized in time slots corresponding to the transmission time of one (maximum sized) data packet plus the necessary tuning delay at each node. Control packets are assumed to be short and the fibres interconnecting the end nodes and the PP in a SAN do typically not exceed a few tens of meters, keeping propagation delays at a fraction of a time slot. One time slot is therefore assumed to provide enough time to send control traffic to the PP, for the PP to run the scheduling algorithm, and to send control messages back to the end nodes. At the end of each time slot, all end nodes tune their respective transmitters and receivers according to the information in the control packets from the PP. Data transmissions take place in the following time slot.

After sorting all packets queued for transmission, each end node chooses the first packet from its HRT queue, or, in case no HRT traffic is present, the first packet from the SRT queue, and sends a control message to the PP. This information specifies source node, destination node and deadline of the packet in question, as well as its traffic class (HRT, SRT or NRT). After sorting all requests according to EDF, the PP checks the control messages one by one, starting with the request with the shortest deadline, and determines whether the requested transmission can be accepted or not. As soon as a destination is booked for the next time slot, any other request for communication with this destination is denied. Accordingly, only one transmission from any specific sending node is allowed per time slot. The PP then simultaneously sends out individual control messages to inform each node whether its sending request was approved and whether it is scheduled to receive data traffic in the following time slot. There is only one single dedicated wavelength connecting each pair of nodes, so a node knows which wavelength to tune its transmitter and receiver to as soon as the identity of its communication partner is known. Nodes that are not scheduled to send or receive data get an empty control packet from the PP for synchronization reasons.

While scheduled data transmissions take place in the next time slot, new control data are exchanged between PP and the end nodes. The parallel handling of data and control traffic is enabled by the individual, fixed wavelength of the control channel from each node and does not interfere with data traffic sent on any of the remaining N–1 wavelengths. Furthermore, the characteristic properties of the AWG allow simultaneous control traffic from all the nodes to be sent to the PP. Therefore, the medium access control results in no more than one single time slot of delay, and does not have any further impact on the data traffic flow (Fig. 5).

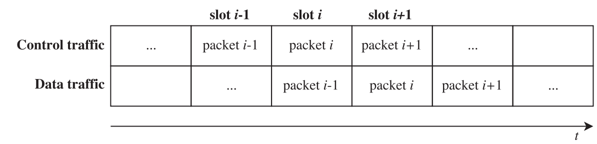

Fig. 5. Control and data traffic flow.

Synchronization in the proposed AWG network is straight-forward. As all the fibres between the nodes and the AWG are assumed to be of equal length, the nodes are synchronized on incoming control packets from the PP. Each node has its own control channel, so clock synchronization on the packet level is not required. Interconnection distances in the targeted networks are assumed to be short, so the difference in propagation delay due to dispersion is negligible. (For a discussion of timing and dispersion in WDM star networks, see [47]).

3.2.2. Prioritization

The network and its belonging MAC protocol support the three indicated traffic classes, however, those traffic classes should not be set equal to priorities. Hard real-time merely indicates that the traffic needs to receive a deadline guarantee in order to be accepted into the network, so that a timely treatment of the packets can be made certain. Soft real-time and non-real-time do not have these kinds of demands. While all HRT traffic is treated before all SRT, which in its turn is treated before all NRT, the actual prioritization mechanism within those traffic classes is deadline-dependent as packets are ordered according to deadline by means of the EDF algorithm.

Obviously this type of traffic scheduling introduces a starvation risk for SRT and NRT traffic, but the target systems specified for this network and protocol solution are systems in need of guaranteed HRT support, and support for SRT or NRT is merely granted in case of available bandwidth. For a system that needs to be able to send a certain amount of SRT/NRT traffic alternative solutions are available, see, e.g., [8,9], or one could use the system suggested in [4], but reserve a certain amount of bandwidth for non-HRT services.

3.3. Simple timing analysis

Assuming, e.g., a 64×64 AWG, the length of a control packet to the protocol processor is less than 100 bits: 6 bits for the source node ID, 6 bits for the destination node ID, 20 bits to state the deadline, 1 bit indicating the traffic class (HRT or SRT/NRT), and some additional bits of overhead. Under the assumption of a bit rate of 2.5 Gb/s, this results in a transmission time, T_{p1}, of 40 ns. With 20 ns, the transmission time, T_{p2}, of a control message from the PP to an end node is even shorter, comprising one flag bit, to accept or deny the requested transmission, and 6 bits to identify the node from which it will receive traffic in the next time slot. Both the processing time, T_r, at the PP and the tuning time, T_t, at the end nodes are assumed to be 100 ns [48]. When defining the length of one time slot, T_s, to be 1 s, this includes time for tuning, leaving 900 ns communication time, i.e., 2250 bits as the maximum data packet length. This results in the following propagation time, T_a:

$$T_a = \frac{T_s - T_{p1} - T_r - T_{p2}}{2} = 420\ ns \tag{1}$$

In other words, an assumed time slot of 1 s leaves enough communication time to support a fibre length of 84 m, assuming a light propagation speed through fibre of $2\cdot 10^8$ m/s, between the PP and any end node (Fig. 6).

3.4. Summary of the case study

In order to verify the feasibility of the proposed network and

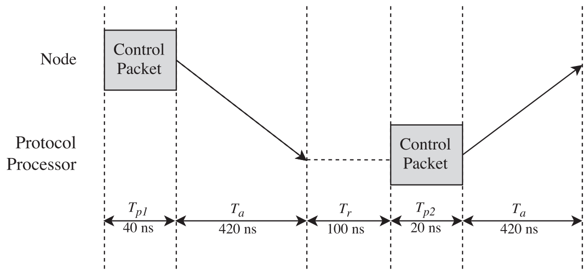

Fig. 6. Control packet propagation.

protocol, the system was evaluated in [4] according to specific application requirements. Radar Signal Processing (RSP) is an application area requiring SANs with support for heterogeneous real-time services. A simulation using an RSP case was carried out and delay, throughput and deadline miss ratio for the network were analysed.

3.4.1. Case definition

In [49], a full case definition for a RSP case scenario is described. The evaluation was inspired by the straight pipeline case, one of three described alternatives. Due to its clear traffic pattern, this case simplifies analysis and understanding of the simulation results, while still requiring three traffic classes.

For the simulation of the RSP system a network architecture with a 16×16 AWG at its centre, consisting of one PP and fifteen communicating end nodes, was chosen. One end node serves as a master node, while the remaining fourteen slave nodes are used for the pipelined data flow (Fig. 7). Further assumptions were periodic traffic, a propagation speed of 2.5 Gb/s and a time slot length of 1 s. In the context of this simulation, the delay caused by tuning was set to be negligible, so the length of one time slot corresponds to one data packet of 2500 bits. The traffic in the RSP case definition consists of three main traffic types: control traffic, which flows in two directions between the master node and each slave node in the straight pipeline, data traffic from node to node in a pipelined manner, and other traffic, such as logging of data or long term statistics from the pipeline nodes to the master node. These three traffic types and their individual deadline requirements are represented by the three traffic classes from the protocol: HRT, SRT and NRT traffic. The relative delay bounds for the real-time traffic classes are equal to their period times.

3.4.2. Case study results

Throughput, average delay and deadline miss ratio of the traffic classes were analysed by using both fixed, predefined traffic parameters (Table 1), and by varying the amount of SRT or HRT traffic in the network.

While the simulation showed good performance results for HRT traffic, the analysis still only could provide a guarantee for a maximum throughput of 1 packet per time slot, namely the packet with the shortest deadline. The parallel nature of the multiwavelength network was not taken into consideration.

4. Improved hard real-time traffic support

As mentioned earlier, HRT support requires the guarantee of timely delivery of packets with strict timing constraints. In order to determine

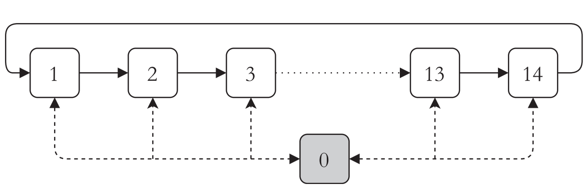

Fig. 7. Master node and straight pipeline of slave nodes according to RSP case definition.

Table 1
Traffic type parameters.

Traffic type	**Traffic class**	**Communication pattern**	**Period [slots]**	**Payload [packets]**	**Delay bound [slots]**
Control	HRT	Master/Slave	100	1	100
Data	SRT	Straight pipeline	5000	4000	5000
Other	NRT	Many-to-one	5000	950	None

the performance characteristics of HRT traffic, a deterministic delay analysis of the proposed network architecture is needed. At any given time a total number of Q real-time traffic flows is present in the network, where each flow q is characterized by a source S_q, destination D_q, period P_q (minimum message interarrival time), deadline E_q, and capacity C_q (maximum message length, expressed as the total transmission time of all packets in the message), and where $1 \le q \le Q$. In order to increase the number of guaranteed packets which can be transmitted concurrently over the single-hop network, traffic dependencies are analysed on a per flow level. The goal of this method is to determine which packets can be sent concurrently in case, e.g., not all nodes request the same destination node.

4.1. Basic real-time analysis

The task of analysing these traffic flows over the network, which resources they use, and if a schedulable traffic allocation is possible, can be mapped onto the problem of uniprocessor task scheduling. In this mapping, the traffic flows are treated as synchronous, periodic tasks that have to be scheduled on the network. The capacity C_q of a traffic flow corresponds to the worst-case execution time (WCET) of the task to be scheduled.

Assuming a worst-case scenario, (i.e., all of the nodes in the network want to send to the same destination in any given time slot), there is always at least one packet, namely the one with the earliest deadline, that can be guaranteed access to the medium due to the assumption of EDF scheduling in the end nodes and the PP. This means that a capacity of one packet per slot always can be utilized by possible HRT traffic. Assuming periodic traffic, where each traffic flow has a deadline equal to its period, basic EDF theory [46] can be used to analyse the system. The first part of the analysis is a utilization check and the utilization of a hard real-time system is defined as:

$$U(q) = \sum_{q=1}^{Q} \frac{C_q}{P_q} \le U_{max} \tag{2}$$

where U is the network utilization by the periodic traffic, C_q is the amount of data per period, P_q is the period of the traffic flow and U_{max} the maximum utilization of the network by HRT traffic that must not be exceeded.

When relaxing the demand of the deadlines being equal to the periods, the analysis must be extended to comprise a second part. We start by introducing the following concepts:

- The hyperperiod (H) is the least common multiple of all periods of the traffic flows, i.e. the length of the time interval from a common starting point of all flows' periods to the point of the next common starting point.
- The busyperiod (B) is any interval within the H during which the link is not idle.
- The workload function $h(t, q)$ is calculated as the sum of the maximum message lengths per period, C_q, of all instances of traffic flows q with an absolute deadline less than or equal to a point in time t, where t signifies the number of time units elapsed since the beginning of the H [50–52], i.e.

$$h(t, q) = \sum_{q=1, E_q \le t}^{Q} \left\lceil \frac{t + P_q - E_q}{P_q} \right\rceil \cdot C_q. \tag{3}$$

In a communication context, a feasibility analysis confirms that the admission of an additional real-time traffic flow still results in a feasible traffic flow set. A feasible real-time traffic flow set in its turn is a traffic flow set where no packet misses its deadline. The feasibility testing is performed in two steps, where the following two constraints have to be met:

Constraint 1 – A set of flows is only allocatable over the network link, if the link utilization is less than or equal to 1 (100%).

Following Eq. (2), this gives

$$U(q) = \sum_{q=1}^{Q} \frac{C_q}{P_q} \le 1. \tag{4}$$

This condition is necessary, but not sufficient to be able to ensure a 100% success rate for the transmission of real-time traffic with guaranteed deadlines, except for the case of all packet deadlines being equal or longer than the period of the traffic flow to which they pertain.

Constraint 2 – For all values of t, the workload function $h(t, q)$ has to be less than or equal to t, i.e.,

$$h(t, q) \le t \quad \forall\, t. \tag{5}$$

This second condition, introduced in [50,51], and generalized in [52], was added to ensure the continued feasibility of the system when adding a new traffic flow.

As this feasibility analysis was developed for real-time systems, it assumes fully preemptive tasks. In network communication, packets normally cannot be preempted and therefore the possibility of further delay has to be taken into account. A constant T_b is defined which denotes the maximum blocking time that the transmission of one maximum sized packet can introduce into the system. Furthermore, the exchange of control information adds an additional delay, T_c. In accordance with earlier reasoning, T_b and T_c will be equally long. The shortened delay bound $E_q{}'$ is:

$$E_q{}' = E_q - T_b - T_c. \tag{6}$$

This means that the workload function is remodelled as follows:

$$h(t, q) = \sum_{q=1, E_q{}' \le t}^{Q} \left\lceil \frac{t + P_q - E_q{}'}{P_q} \right\rceil \cdot C_q. \tag{7}$$

In order to decrease the high computational complexity, [53] reduce the number of necessary instances of evaluation to the instances of t where

$$t \in \bigcup_q \{m \cdot P_q + E_q \colon m = 0, 1, 2, \ldots\}, \quad t \in [1; B_1] \tag{8}$$

and where B_1 is the first B in the first H of the schedule where all periods start at time zero. For non-preemptive communication, the delay bound has to be further adjusted according to Eq. (6), which gives:

$$t \in \bigcup_q \{m \cdot P_q + E_q{}' \colon m = 0, 1, 2, \ldots\}, \quad t \in [1; B_1]. \tag{9}$$

Only when both the utilization constraint and the workload constraint are fulfilled can a feasible traffic allocation be guaranteed.

4.2. Traffic analysis algorithm

Applying the specified feasibility analysis to a set of real-time traffic flows over a network will return a simple ‘yes/no’ answer, stating if all flows could be allocated over “the network”, which is seen as a single resource. This stems from uniprocessor task scheduling as the original application of this analysis. In reality, a network is a set of overlapping resources, depending upon the physical and logical network architecture and the MAC method used. In the presented context, a resource is constituted by a sender, a receiver and the unidirectional light path connecting them. If any of those components is busy, the whole resource is seen as reserved. This leads to the following definitions. .

Definition 1. A **resource** consists of a sender, a receiver and the unidirectional light path connecting them.

Definition 2. A resource is defined to be **busy** if either the sender at the source node is transmitting to any other node connected to it, or the receiver at the destination is receiving from any other node connected to it, or both are participating in the communication over the unidirectional link between them.

The sender at each source node can only send to one destination at a time, and the receiver at each destination node can only receive from one source at a time (not including the communication with the protocol processor on a separate control channel). Consequently, the resource, and therefore the unidirectional link between those two nodes, cannot be used for other communication as soon as one of them is busy.

The original feasibility analysis contains a rather large amount of pessimism as it ignores the possibility of simultaneous transmissions over independent resources. In the following, we define virtual overlapping subgroups in order to analyse them individually according to the feasibility test. .

Definition 3. One **subgroup** consists of one main traffic flow and all other traffic flows it shares at least one part of the resource with, i.e., the flows with the same source or destination node (or both) as the main traffic flow.

Definition 4. Two subgroups are **overlapping** if they share at least the sender at the source node or the receiver at the destination node. As long as two subgroups are not overlapping, concurrent transmissions in them can occur at any time.

A subgroup, which is a set of real-time traffic flows, has to be found for each individual traffic flow, and both the utilization and the workload tests have to be applied on all these subgroups. In order to analyse the interdependencies of the traffic flows contained in one subgroup, each single flow has to be checked for shared resources with any other flow. The studied traffic flow F_i is characterized by its source S_i, destination D_i, period P_i, deadline E_i, and capacity C_i. The link between S_i and D_i is denoted L_i. A resource R_i includes S_i, D_i and the link between them, L_i. As a first step, we create a set M_i, i.e., the subgroup with all traffic flows F_j that either have S_i as its source node or D_i as its destination node, i.e.,

$$M_i = \bigcup_{j=1}^{Q} \{F_j | S_j = S_i \;\text{ or }\; D_j = D_i\}. \tag{10}$$

This set also includes the studied flow F_i itself and those flows parallel to it, i.e., flows that share all parts of the resource with the currently studied flow (i.e., it is an inclusive OR not exclusive OR in the set definition).

4.3. Improved, less pessimistic feasibility analysis

Instead of applying the feasibility analysis to all traffic flows (as in Section 4.1), it is now applied to each subgroup to decide if all flows competing for the same resource can be allocated without causing deadline misses. The number of subgroups is equal to the number of traffic flows, Q. Based on the previously introduced traffic interdependency analysis, the utilization U_i of R_i by all flows F_j in M_i is now calculated as

$$U_i = \sum_{j=1}^{Q} e_j \quad \text{where } e_k = \begin{cases} \frac{C_k}{P_k}, & \text{if } F_k \in M_i \\ 0, & \text{otherwise} \end{cases} \tag{11}$$

corresponding to Eq. (4) in the previous, pessimistic version of the analysis.

This being a worst-case analysis, the necessary condition of $U_i \le 1$ must still be fulfilled. The workload function (introduced in Eq. (7)) is now applied to the smaller subgroup M_i, checking if HRT traffic can be guaranteed for the studied flow F_i. The workload is now calculated as

$$h_i(t) = \sum_{j=1}^{Q} g_j$$
where
$$g_k(t) = \begin{cases} \left\lfloor \frac{t + P_k - E_k'}{P_k} \right\rfloor \cdot C_q, & \text{if } F_k \in M_i \text{ and } E_k' \le t \\ 0, & \text{otherwise.} \end{cases} \tag{12}$$

A negative answer does in itself not mean that the flow automatically will miss its deadline, but simply that no guarantee for it can be given. Both utilization and workload check are carried out for each set of traffic flows, i.e., each single traffic flow will be studied as the main flow of a set.

Let X denote the set containing all traffic flows for which timely treatment can be guaranteed, and Y be the size of this set, i.e., the actual number of flows, for which deadline guarantees can be provided, then

$$X = \bigcup_{j=1}^{Q} \{M_i | U_i \le 1 \text{ and } h_i(t) \le t, \quad \forall\, t\} \tag{13}$$

and

$$Y = |X|. \tag{14}$$

When analysing a hard real-time system, Y must be equal to the number of traffic flows in the system, i.e., Q, for the system to be feasible. However, for the analysis of soft real-time systems, Y can be used as a parameter indicating the degree to which real-time traffic can be guaranteed.

There can still be pessimism contained in this analysis. Flows might be included in several subgroups, but due to the possibility of intricate traffic interdependencies, the complexity of an analysis studying dependencies more than one step from the main studied traffic flow might grow immensely fast, and is, for now, outside the scope of this work. However, the presented approach can increase the capacity that can be guaranteed for HRT traffic from a throughput of one packet per time slot to the throughput reached by Y traffic flows instead.

5. Evaluation of the real-time scheduling performance

In order to demonstrate the effect of the improved feasibility analysis, an extension to the Java-based simulator used in [4] was implemented. Still assuming a 16×16 AWG, resulting in 15 end nodes and the protocol processor, the traffic pattern was assumed to be the following. The period as well as the deadline of all traffic flows is 100 time slots and their maximum message length corresponds to the length of one time slot. The source of each traffic flow is randomized with an even distribution, while the choice of destination is limited to include nodes contained in destination groups of a certain (variable) size. The maximum number of channels requested in the system is set to 2000. Each data point in the evaluation graphs is the result of 100 iterations in order to increase the statistical reliability of the result.

In Fig. 8, the number of requested traffic flows is increased (in steps of one) from 1 to 2000, after which the network seems to be saturated by the traffic load for all curves, i.e., for all destination group sizes. The figure shows the throughput in packets per time slot experienced by the guaranteed traffic flows versus the number of requested traffic flows. (The calculation of the actual number of accepted traffic flows is basic as each channel q has a bandwidth utilization of 1% (C_q=1 slot, P_q=100 slots)). The graphs illustrate the case when each source can send to 1, 4, 7 or a maximum of 14 possible other destinations. (All possible sizes of destination groups were simulated, but some are excluded from the figure for readability reasons.) The destinations included in each destination group are randomized with an even distribution.

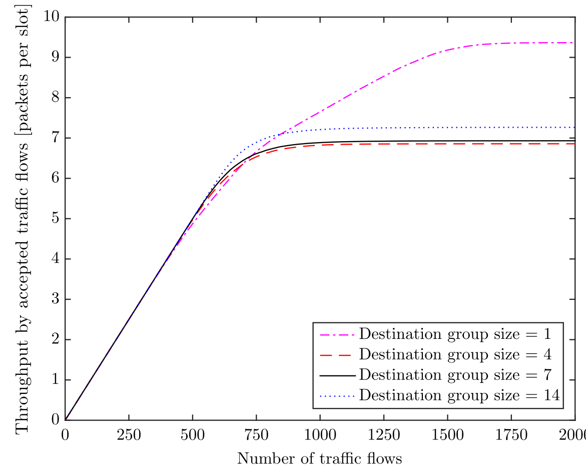

Fig. 8. Throughput by guaranteed traffic flows.

As shown in Fig. 8, the maximum throughput by accepted traffic flows is reached at the network saturation point of about 800 requested traffic flows. Most prominent, however, is the graph for a destination group size of one, where the maximum throughput is 9.53, instead of the remaining values of around 7. The reason for this deviating behaviour is the relatively low probability of overlapping destination groups when each traffic flow from a particular source has to have the same destination.

The theoretical throughput maximum, denoted Z, reachable in a network with this configuration, and under the assumption of equally many traffic flows between any source-destination pair, can be calculated by

$$Z = \frac{N \cdot N_d}{2N_d - 1} \tag{15}$$

where N is the total number of nodes in the network, and N_d is the number of destinations per destination group. This theoretical average throughput should however only be used for an approximate comparison since it is based on the assumption of a non-random even distribution. The theoretically calculated values are compared to the simulation results in Fig. 9, where the difference between the simulated and calculated throughput decreases continuously as the number of possible destinations per source increases. For a small value of N_d, the lower simulated throughput is due to the relatively high probability of non-overlapping destination groups, which reduces the actual throughput compared to the calculated number. The turning point where the simulated throughput changes from a sinking to an increasing trend lies around a destination group size of four.

The throughput of the guaranteed traffic flows versus N_d is studied in Fig. 10, showing the results for different traffic loads in the system. While 500 traffic flows can be easily accommodated by the network, higher traffic loads experience the same behaviour as seen in Fig. 9. However, this graph also shows the effect of a low probability of overlapping destination groups (which was illustrated by the ‘Destination group size=1’ graph in Fig. 8 earlier), observable up to a destination group size of four different destinations per source.

The simulation results verify that the number of traffic flows that

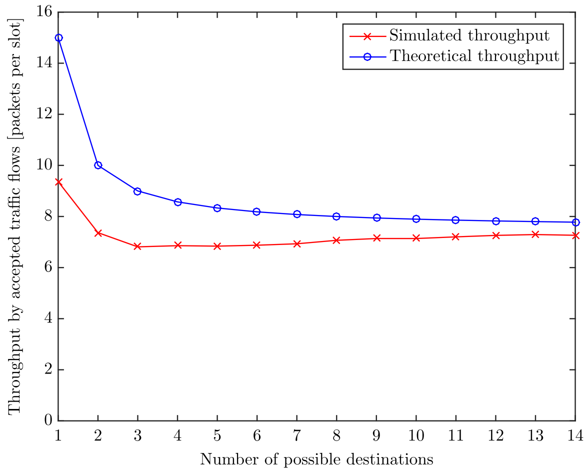

Fig. 9. Calculated vs. simulated throughput.

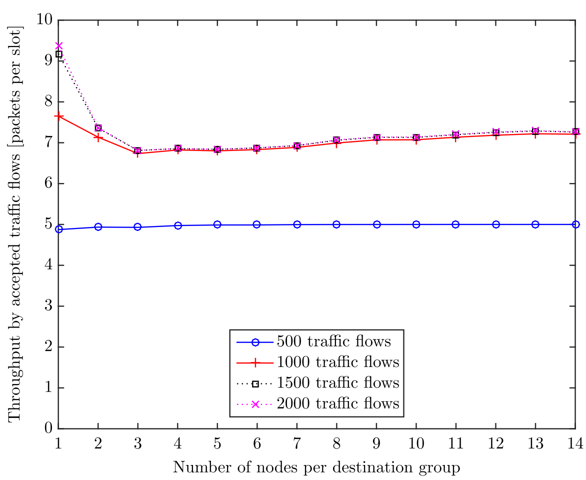

Fig. 10. Throughput depending on destination group size.

can be guaranteed to meet their deadlines can be increased considerably by combining the real-time feasibility analysis with the suggested traffic analysis. The guaranteed throughput was increased from one packet per slot to around seven, i.e., a utilization of about 700%. For some traffic patterns, even higher guaranteed throughputs can be achieved, as shown by the simulations.

6. Conclusions

High-performance embedded systems with heterogeneous traffic, high throughput needs and strict delay bound demands are in need of new communication solutions providing the necessary performance. Optical technologies are an interesting contender due to their high bandwidth and low loss qualities. In this paper, special interest is directed towards communication in a AWG-based network as the AWG due to its potential for high concurrency offers the possibility of short communication delays and numerous parallel transmissions.

When targeting the support for hard real-time systems, guaranteed throughput and an upper-bounded delay are vital. Employing a deterministic medium access protocol combined with an admission control method including a schedulability analysis can provide the possibility of hard real-time guarantees. However, the real-time analysis method, originating from uniprocessor scheduling and presented in earlier work to achieve timing guarantees, did not take into consideration the inherent parallelism in multiwavelength optical networks. Our work adapted the existing analysis to a multichannel communication context by adding a traffic dependency analysis and thereby addressing the pessimism included in the earlier analysis.

Evaluation of the new solution shows an considerable improvement of the performance of the analysis for multichannel networks. A simulation of the analysis for different traffic patterns, indicates that the throughput guarantee can be improved substantially. Based on an example of a radar signal processing application and a network architecture with a 16×16 AWG at its core, we find an improvement of the guaranteed traffic throughput of around 700%. Thus, with our new analysis, we can accept a significantly higher amount of hard real-time traffic into the network, while still being able to provide guarantees for real-time performance. This creates the possibility of catering for the needs of more demanding real-time applications without changing the existing AWG-based network architecture or the utilized deterministic MAC protocol.

In this paper we have shown that analysing traffic dependencies is a valid and relatively straight forward approach for increasing the amount of guaranteed real-time traffic in this type of wavelength-routed star networks. However, even other types of multichannel networks with parallel transmissions can benefit of this method. Open questions include the adaptation of this method to, e.g., ring or mesh networks. Even wireless multifrequency networks are a possible application area. So far, this method has merely been applied to single-hop networks, as traffic interdependencies along a multihop path will increase in complexity very rapidly. However, assuming static routing paths, traffic analysis should be possible even for multihop networks.

References

[1] J. Cao, A.T. Chan, Y. Sun, S.K. Das, M. Guo, A taxonomy of application scheduling tools for high performance cluster computing, Clust. Comput. 9 (3) (2006) 355–371.
[2] S. Haykin, Cognitive radar: a way of the future, IEEE Signal Process. Mag. 23 (1) (2006) 30–40.
[3] N. Matveeva, Y. Sheynin, E. Suvorova, QoS support in embedded networks and NoC, in: Proceedings of the 16th Conference of Open Innovations Association (FRUCT16), IEEE, 2014, pp. 51–59.
[4] A. Böhm, M. Jonsson, K. Kunert, H. Hoang, A fibre-optic AWG-based real-time network and its applicability to high-performance embedded computing, in: Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium, 2005, pp. 133b–133b.
[5] H. Hoang, M. Jonsson, Switched real-time ethernet in industrial applications-deadline partitioning, in: Proceedings of the 9th Asia-Pacific Conference on Communications (APCC 2003), vol. 1, 2003, pp. 76–81.
[6] B. Mukherjee, WDM-based local lightwave networks. I. single-hop systems, IEEE Netw. 6 (3) (1992) 12–27.
[7] M.K. Smit, C. Van Dam, PHASAR-based WDM-devices: principles, design and applications, IEEE J. Sel. Top. Quantum Electron. 2 (2) (1996) 236–250.
[8] M.S. Borella, J.P. Jue, B. Mukherjee, Simple scheduling algorithms for use with a waveguide grating multiplexer based local optical network, Photon. Netw. Commun. 1 (1) (1999) 35–48.
[9] K. Bengi, Performance of single-hop WDM lans supporting real-time services, Photon. Netw. Commun. 1 (4) (1999) 287–301.
[10] M.J. Spencer, M.A. Summerfield, WRAP: a medium access control protocol for wavelength-routed passive optical networks, J. Lightwave Technol. 18 (12) (2000) 1657–1676.
[11] C. Fan, S. Adams, M. Reisslein, The FTΛ-FRΛ AWG network: a practicalsingle-hop metro WDM network for efficient uni-and multicasting, J. Lightwave Technol. 23 (3) (2005) 937–954.
[12] M. Maier, M. Reisslein, A. Wolisz, A hybrid MAC protocol for a metro WDM network using multiple free spectral ranges of an arrayed-waveguide grating, Comput. Netw. 41 (4) (2003) 407–433.
[13] T. Maier, A. Wolisz, Single-hop WDM network with high spectrum reuse based on an arrayed waveguide grating, in: Proceedings of the IEEE/ACM/SPIE Optical Netw. Workshop, 2000.
[14] M. Maier, M. Reisslein, A. Wolisz, Towards efficient packet switching metro WDM networks, Opt. Netw. Mag. 3 (6) (2002) 44–62.
[15] H.-S. Yang, M. Herzog, M. Maier, M. Reisslein, Metro WDM networks: performance comparison of slotted ring and AWG star networks, IEEE J. Sel. Areas Commun. 22 (8) (2004) 1460–1473.
[16] S.A.M. Kirmani, H.-S. Yang, S. Shin, A wavelength reusable metro WDM star network with a hybrid MAC protocol, in: Proceedings of the Conference on Lasers and Electro-Optics - Pacific Rim (CLEO/Pacific Rim '07), Optical Society of America, 2007, pp. 1–2.
[17] H.-T. Lin, C.-L. Lai, T.-S. Wang, Y.-C. Huang, Supporting triple-play services with private networking over WDM EPONs, in: 2014 Proceedings of the 14th International Symposium on Communications and Information Technologies (ISCIT), 2014, pp. 414–418.

[18] T. Uesugi, R. Kubo, Dynamic wavelength allocation technique with multicast-capable awg router for energy-efficient intra-datacenter networks, in: 2015 Proceedings of the 20th Microoptics Conference (MOC), 2015, pp. 1–2.
[19] H. Rastegarfar, L.A. Rusch, A. Leon-Garcia, Optical load-balancing tradeoffs in wavelength-routing cloud data centers, J. Opt. Commun. Netw. 7 (4) (2015) 286–300.
[20] H. Kim, B. Shin, J. Lee, C. Un, Performance evaluation of reservation protocol with priority control for single-hop WDM networks, Electron. Lett. 31 (17) (1995) 1472–1473.
[21] M. Jonsson, K. Nilsson, B. Svensson, A fiber-optic interconnection concept for scaleable massively parallel computing, in: Proceedings of the Second International Conference on Massively Parallel Processing Using Optical Interconnections (MPPOI'95), 1995, pp. 313–320.
[22] M. Jonsson, K. Börjesson, M. Legardt, Dynamic time-deterministic traffic in a fiber-optic WDM star network, in: Proceedings of the Ninth Euromicro Workshop on Real-Time Systems (ECRTS'97), 1997, pp. 25–33.
[23] H.-Y. Tyan, J.C. Hou, B. Wang, C.-C. Han, On supporting temporal quality of service in WDMA-based star-coupled optical networks, IEEE Trans. Comput. 50 (3) (2001) 197–214.
[24] L. Dong, R. Melhem, D. Mossé, Time slot allocation for real-time messages with negotiable distance constraints, in: Proceedings of the Fourth IEEE Real-Time Technology and Applications Symposium, 1998, pp. 131–136.
[25] B. Li, Y. Qin, Traffic scheduling in a photonic packet switching system with QoS guarantee, J. Lightwave Technol. 16 (12) (1998) 2281–2295.
[26] A.C. Kam, K.-Y. Siu, R.A. Barry, E.A. Swanson, Toward best-effort services over WDM networks with fair access and minimum bandwidth guarantee, IEEE J. Sel. Areas Commun. 16 (7) (1998) 1024–1039.
[27] K. Bengi, Medium access control protocols for passive-star WDM single-hop networks efficiently integrating real-time and data services, Opt. Netw. Mag. 2 (4) (2001) 88–100.
[28] M. Ma, M. Hamdi, Providing deterministic quality-of-service guarantees on WDM optical networks, IEEE J. Sel. Areas Commun. 18 (10) (2000) 2072–2083.
[29] X. Huang, M. Ma, Y. Ma, Z. Lin, An efficient QoS prediction scheme in single-hop passive star coupled WDM optical networks, in: IFIP International Conference on Wireless and Optical Communications Networks (WOCN'07), 2007, pp. 1–5.
[30] X. Huang, M. Ma, A heuristic adaptive qos prediction scheme in single-hop passive star coupled WDM optical networks, J. Netw. Comput. Appl. 34 (2) (2011) 581–588.
[31] P.G. Sarigiannidis, Applying quality of service prediction in WDM optical networks, in: Proceedings of the 15th IEEE Mediterranean Electrotechnical Conference (MELECON 2010), 2010, pp. 1406–1410.
[32] M. Bokhari, P. Saengudomlert, Analysis of mean packet delay for upstream transmissions in passive optical networks with sleep mode, Opt. Switch. Netw. 10 (3) (2013) 195–210.
[33] S. Parkes, C. McClements, D. McLaren, A.M. Martinez, A.F. Florit, A.G. Villafranca, SpaceFibre: Adaptive high-speed data-link for future spacecraft onboard data handling, in: 2014 NASA/ESA Conference on Adaptive Hardware and Systems (AHS), 2014, pp. 164–171.
[34] N. Matveeva, Y. Sheynin, E. Suvorova, QoS in SpaceFibre and SpaceWire/GigaSpaceWire protocols, in: 2014 International SpaceWire Conference (SpaceWire), 2014, pp. 1–8.
[35] Y. Sheynin, E. Suvorova, N. Matveeva, A. Khakhulin, I. Orlovsky, D. Romanov, SpaceFibre based spacecraft network case study, in: 2014 International SpaceWire Conference (SpaceWire), 2014, pp. 1–5.
[36] I.-F. Chao, S.-T. Chang, A novel access and control system in a WDM passive star-coupled network for optical backhaul, J. Opt. Commun. Netw. 6 (12) (2014) 1082–1092.
[37] J. Diao, P.L. Chu, Packet rescheduling in WDM star networks with real-time service differentiation, J. Lightwave Technol. 19 (12) (2001) 1818–1828.
[38] J.S. Choi, Performance analysis of a bandwidth guaranteed medium access control protocol for WDM local networks, Photon. Netw. Commun. 6 (1) (2003) 69–81.
[39] J.M.M. Bellón, M. Jonsson, EDF-based protocol for a fiber optic network with passive star and separate control channel, in: Proceedings of the IASTED International Conference on Communication and Computer Networks (CCN'04), 2004, pp. 13–18.
[40] L. Wang, M. Hamdi, R. Manivasakan, D.H. Tsang, Multimedia-MAC protocol: its performance analysis and applications for WDM networks, IEEE Trans. Commun. 54 (3) (2006) 518–531.
[41] S. Petridou, P. Sarigiannidis, G. Papadimitriou, A. Pomportsis, Nodes clustering in WDM star networks with real-time traffic, in: Proceedings of the 14th IEEE Mediterranean Electrotechnical Conference (MELECON 2008), 2008, pp. 252–257.
[42] X. Huang, M. Ma, A performance model for differentiated service over single-hop passive star coupled WDM optical networks, J. Netw. Comput. Appl. 34 (1) (2011) 183–193.
[43] R. Deepalakshmi, C.S. Kumar, S. Rajaram, New MAC protocol for traffic routing in optical networks by exploiting delays in dynamic bandwidth allocation, in: Proceedings of the 2011 IEEE International Conference on Computer Applications and Industrial Electronics (ICCAIE), 2011, pp. 325–330.
[44] M.R. Radivojević, P.S. Matavulj, Highly flexible and efficient model for QoS provisioning in WDM EPON, J. Opt. Commun. Netw. 5 (8) (2013) 921–931.
[45] V.V. Mai, A.T. Pham, Integrated FSO/PON for broadband access networks: a comprehensive protocol stack design and analysis, in: Proceedings of the 2015 IEEE Global Communications Conference (GLOBECOM'15), 2015, pp. 1–7.
[46] C.L. Liu, J.W. Layland, Scheduling algorithms for multiprogramming in a hard-real-time environment, J. ACM 20 (1) (1973) 46–61.
[47] G. Semann, P. Humblet, Timing and dispersion in WDM optical star networks, in: Proceedings of the Twelfth Annual Joint Conference of the IEEE Computer and Communications Societies. Networking: Foundation for the Future (INFOCOM'93), 1993, pp. 573–577.
[48] R. Ramaswamy, N. Weng, T. Wolf, Characterizing network processing delay, in: Proceedings of the 2004 IEEE Global Telecommunications Conference, (GLOBECOM'04), vol. 3, 2004, pp. 1629–1634.
[49] C. Bergenhem, M. Jonsson, B. Gördén, A. Åhlander, Heterogeneous real-time services in high-performance system area networks-application demands and case study definitions, Tech. Rep. IDE Tech. Rep. 263, Halmstad University, Halmstad, Sweden, October 2002.
[50] S.K. Baruah, A.K. Mok, L.E. Rosier, Preemptively scheduling hard-real-time sporadic tasks on one processor, in: Proceedings of the 11th Real-Time Systems Symposium (RTSS'90), 1990, pp. 182–190.
[51] S.K. Baruah, L.E. Rosier, R.R. Howell, Algorithms and complexity concerning the preemptive scheduling of periodic, real-time tasks on one processor, Real.-Time Syst. 2 (4) (1990) 301–324.
[52] M. Spuri, Analysis of deadline scheduled real-time systems, Tech. Rep. Tech. rep. no. 2772, INRIA, France, January 1996.
[53] J.A. Stankovic, M. Spuri, K. Ramamritham, G.C. Buttazzo, Deadline scheduling for real-time systems: EDF and related algorithms, vol. 460, Springer Science & Business Media, 2012.